# **WEST Search History**

emed ebith Restore Cencel

DATE: Thursday, February 23, 2006

Hide?	<u>Set</u> Name	Query	<u>Hit</u> Count
	DB=B	PGPB, USPT, USOC, EPAB, JPAB, DWPI, TDBD; PLUR=YES; OP=ADJ	
	L57	L56 and (data near5 sets)	7
	L56	L55 and attributes	7
	L55	L54 and (mining near5 pattern)	9
	L54	mining near5 template	47
	L53	6061682 .uref.	17
	L52	L51 and @py<=2003	11
	L51	(discover\$3 and mining).ti.	20
	L50	(knowledge and mining).ti.L48	0
	L49	(discover\$3 and mining).ti.L48	0
	L48	6292797 .uref.	6
	L47	L46 and attributes	4
	L46	(mining and pattern and discovery).ti,ab. and @py<=2003	17
	L45	(mining and pattern and discovery).ti.	0
	L44	(mining and pattern and discovery).ti,ab.	21
	L43	(pattern near5 join\$3) and (pattern near5 template) and relational and database\$1 and mining and tables	3
	L42	(5809499 or 6434570).pn.	4
	L41	relational and databases and knowledge and discovery and pattern and merge and join and sort\$3 and nodes and templates and attributes and tables and values and preference\$1 and @py<=2003	38
	L40	L39 and algorithms	6
	L39	L38 and (knowledge near5 base)	7
	L38	L37 and mining	39
	L37	L36 and attributes	39
	L36	L35 and mining	49
	L35	merge and joining and discovery	613
	L34	'merge-joining' and discovery	1
	L33	'merge-joining' and data and tables and mining	1
	L32	L30 and (mining near5 preference\$1)	0
	L31	L30 and (mining near5 preferences)	0
	L30	L29 and (user near5 interfac\$3)	24

L29	L28 and (data near5 mining)	24
L28	L27 and (data near5 set\$1)	24
L27	L24 and hierarchical	25
L26	L24 and 'anti-monotonicity'	0
L25	L24 and monotonicity	0
L24	L23 and templates	25
L23	knowledge and data and mining and join and merge and pattern and attributes and tables and fields and rules and algorithms and @py<=2003	29
L22	'anti-monotonicity property'	4
L21	antimonotonicity property	0
L20	L19 and templates	2
L19	L18 and discover\$3	4
L18	L17 and pattern	4
L17	L16 and (multiple near5 attributes)	12
L16	relational near5 dataset\$1	58
L15	110 and pattern	13
L14	112 and join	1
L13	L12 and merge	1
L12	L10 and relationships	10
L11	L10 and mining and templates	1
L10	(discovery near5 techniques) and (multiple near5 attributes)	18
L9	L8 and (discover\$3 near5 (attribute\$1 or field\$1))	5
L8	L7 and template\$1	20
L7	L6 and (mining near5 rule\$1)	135
L6	mining near5 (attribute\$1 or field\$1)	1706
L5	(multiple near5 data) and (data near5 sets) and (data near5 mining) and (mining near5 templates) and ((discover\$3 or knowledge) near5 (attribute\$1 or table\$1))	4
L4	(multiple near5 data) and (data near5 sets) and (data near5 mining) and (mining near5 templates) and ((discover\$3 or knowledge) near5 (attribute\$1 or table\$1)) and @py<=2003	0
L3	(data near5 mining) and (anti\$monotonicity)	5
L2	L1 and (anti\$monotonicity)	0
L1	(data and mining and attribute\$1 and data and set\$1 and relational and (multiple near5 item\$1) and template\$1) and @py<=2004	33

# END OF SEARCH HISTORY

⊠e-mail



Home | Login | Logout | Access Information | Alerts |

#### Welcome United States Patent and Trademark Office

	Sea	rch	Resu	its
--	-----	-----	------	-----

**BROWSE** 

Results for "( ( mining<in>metadata ) <and> ( template<in>metadata ) )<and> ( pattern<..."

SEARCH

IEEE XPLORE GUIDE

Your search matched 8 of 1320520 documents. A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order. » Search Options **Modify Search** View Session History ( ( mining<in>metadata ) <and> ( template<in>metadata ) )<and> ( pattern<in>metad Search **New Search** Check to search only within this results set » Key IEEE Journal or IEEE JNL Magazine ∠ view selected items Select All Deselect All **IEE JNL** IEE Journal or Magazine IEEE Conference **IEEE CNF** 1. Mining generalized query patterns from web logs Proceeding Ling, C.X.; Jianfeng Gao; Huajie Zhang; Weining Qian; Hongjiang Zhang; **IEE Conference IEE CNF** System Sciences, 2001. Proceedings of the 34th Annual Hawaii International ( Proceeding Jan 3-6 2001 Page(s):9 pp. IEEE STD IEEE Standard AbstractPlus | Full Text: PDF(136 KB) IEEE CNF Rights and Permissions 2. An evolutionary approach to pattern-based time series segmentation П Fu-Lai Chung; Tak-Chung Fu; Ng, V.; Luk, R.W.P.; Evolutionary Computation, IEEE Transactions on Volume 8, Issue 5, Oct. 2004 Page(s):471 - 489 Digital Object Identifier 10.1109/TEVC.2004.832863 AbstractPlus | References | Full Text: PDF(1312 KB) | IEEE JNL Rights and Permissions 3. Evolutionary time series segmentation for stock data mining 

4. Wavelet based UXO detection

Rights and Permissions

Hodgson, S.; Dunstan, N.; Murison, R.;

Fu-lai Chung; Tak-chung Fu; Luk, R.; Ng, V.;

Data Mining, 2002. ICDM 2002. Proceedings. 2002 IEEE International Confere

Data Mining, 2002. ICDM 2002. Proceedings. 2002 IEEE International Confere

9-12 Dec. 2002 Page(s):617 - 620

9-12 Dec. 2002 Page(s):83 - 90

Digital Object Identifier 10.1109/ICDM.2002.1184012

Digital Object Identifier 10.1109/ICDM.2002.1183889

<u>AbstractPlus</u> | Full Text: <u>PDF</u>(600 KB) IEEE CNF

AbstractPlus | Full Text: PDF(288 KB) | IEEE CNF

Rights and Permissions

5. A semantic-based image retrieval system: VisEngine

Jian-Yong Sun; Zheng-Xing Sun; Ruo-Hong Zhou; Hui-Feng Wang; Machine Learning and Cybernetics, 2002. Proceedings. 2002 International Con

Volume 1, 4-5 Nov. 2002 Page(s):349 - 353 vol.1 Digital Object Identifier 10.1109/ICMLC.2002.1176772

AbstractPlus | Full Text: PDF(411 KB) | IEEE CNF

#### Rights and Permissions

6.	Damage pattern mining in hurricane image databases Shu-Ching Chen; Mei-Ling Shyu; Chengcui Zhang; Tang, W.Z.; Zhang, K.; Information Reuse and Integration, 2003. IRI 2003. IEEE International Confere 2003 Page(s):227 - 234 Digital Object Identifier 10.1109/IRI.2003.1251418
	AbstractPlus   Full Text: PDF(603 KB) IEEE CNF Rights and Permissions
7.	Disease-specific probabilistic brain atlases Thompson, P.; Mega, M.S.; Toga, A.W.; Mathematical Methods in Biomedical Image Analysis, 2000. Proceedings. IEE 11-12 June 2000 Page(s):227 - 234 Digital Object Identifier 10.1109/MMBIA.2000.852382
	AbstractPlus   Full Text: PDF(556 KB) IEEE CNF Rights and Permissions
8.	HierarchyScan: a hierarchical similarity search algorithm for databases of sequences Chung-Sheng Li; Yu, P.S.; Castelli, V.; Data Engineering, 1996. Proceedings of the Twelfth International Conference of 26 Feb1 March 1996 Page(s):546 - 553 Digital Object Identifier 10.1109/ICDE.1996.492205  AbstractPlus   Full Text: PDF(720 KB)   IEEE CNF
	Rights and Permissions

Indexed by Inspec

Help Contact Us Privacy &: © Copyright 2006 IEEE -



Home | Login | Logout | Access Information | Alerts |

#### Welcome United States Patent and Trademark Office

Search Res	ults		E	BROWSE	SEARCH	IEE	E XPLORE GUIL	DE
Your search	"( ( mining <in>metadata ) &lt; n matched 2 of 1320520 doc. n of 100 results are displayed</in>	ıments.	•		•			<b>⊠</b> e-mail
» Search O	otions							
View Session	on History	Modify	Search			:		
New Search	1	( ( mining <in>metadata ) <and> ( template<in>metadata ) )<and> ( attributes<in>metadata ) Search</in></and></in></and></in>						
		☐ Che	eck to sear	ch only within th	his results set			
» Key		Display	Format:	Citation	C Citation 8	Abstract		
IEEE JNL	IEEE Journal or Magazine							
IEE JNL	IEE Journal or Magazine	t view :	selected it	ems Select	All Deselect	All		
IEEE CNF	IEEE Conference Proceeding IEE Conference	□ 1	. <b>Trends i</b> Wijsen, .	in databases: ı J.;	reasoning and	mining		
	Proceeding			ge and Data Er 13, Issue 3, M				
IEEE STD	IEEE Standard			bject Identifier			430	
				Plus   Reference nd Permissions		<u>PDF(</u> 276 KB)	IEEE JNL	
		☐ <sup>2</sup>	Ke Wang Data Min 27-30 No Digital O Abstract	te-Based Priva g; Fung, B.C.M. hing, Fifth IEEE bv. 2005 Page(s bject Identifier Plus   Full Text: nd Permissions	; Yu, P.S.; International C s):466 - 473 10.1109/ICDM.2 PDF(272 KB)	onference on 2005.142	ation Problems	3

Indexed by #Inspec Help Contact Us Privacy &: © Copyright 2006 IEEE -



Home | Login | Logout | Access Information | Alerts |

	RELEASE 2.1	Welcome United States Patent and Trademark Office					
Search Res	sults	•	BROW	SE SEARC	н	IEEE XPLORE GUIDE	
Your searc	"( ( join <in>metadata ) <ar h matched 5 of 1320520 doo n of 100 results are displayed</ar </in>	cuments.				<b>⊠</b> е-таіl	
» Search O	ptions						
View Sessi	on History	Modif	y Search				
New Searc	<u>h</u>	( ( join•	in>metadata ) <and< td=""><td>&gt; ( merge<in>metadata</in></td><td>))<and>(</and></td><td>attributes<in>metadata Search</in></td></and<>	> ( merge <in>metadata</in>	)) <and>(</and>	attributes <in>metadata Search</in>	
		☐ ci	neck to search only	within this results se	et :		
» Key		Displa	y Format: 📵	Citation C Citation	n & Abstra	ct	
IEEE JNL	IEEE Journal or Magazine						
IEE JNL	IEE Journal or Magazine	t view	selected items	Select All Desele	ct All		
IEEE CNF	IEEE Conference Proceeding	П	1. Parallel collec	tion-equi ioin algorit	thms for a	bject-oriented databases	
IEE CNF	IEE Conference Proceeding	I_i	Taniar, D.; Rah	ayu, J.W.;		osium, 1998. Proceedings. IDEA	
IEEE STD			International	Page(s):159 - 168			
				dentifier 10.1109/IDE	AS.1998.6	94376	
			AbstractPlus   F Rights and Per	Full Text: <u>PDF</u> (80 KB) <u>missions</u>	IEEE C	NF	
			Taniar, D.; Rah High Performar 28 April-2 May	ayu, W.;	Information	ject-oriented collection join que on Superhighway, 1997. HPC A:	
				Full Text: PDF(532 KE			
			Liang Jin; Cher Database Syste International Co 26-28 March 20 Digital Object to	onference on 003 Page(s):137 - 146 dentifier 10.1109/DAS	plications, S FAA.2003		
			Rights and Pen	Full Text: <u>PDF</u> (494 KE missions	3) IEEE (	CNF	
			Al-Kateb, M.; M Temporal Repri 23-25 June 200 Digital Object Io	5 Page(s):83 - 90 lentifier 10.1109/TIM	awi, M.E.; oning, 2005 E.2005.16	5. TIME 2005. 12th Internationa	
			AbstractPlus   F Rights and Pen	full Text: <u>PDF</u> (120 KE <u>nissions</u>	B) IEEE (	CNF	
			Baoliang Liu; Ji	for massive data or anzhong Li; Yanqiu Z	hang;	storage system	

<u>International</u>

7-9 July 2004 Page(s):271 - 276
Digital Object Identifier 10.1109/IDEAS.2004.1319800

<u>AbstractPlus</u> | Full Text: <u>PDF(</u>286 KB) IEEE CNF
<u>Rights and Permissions</u>

Indexed by

Help Contact Us Privacy &:

© Copyright 2006 IEEE -



Home | Login | Logout | Access Information | Alerts |

#### Welcome United States Patent and Trademark Office

☐ Search	Results	
----------	---------	--

Search Res	sults		BROV	VSE	SEARCH	: 18	EEE XPLORE GL	IIDE
Your search	"( ( join <in>metadata ) <ano n matched 5 of 1320520 doc n of 100 results are displayed</ano </in>	uments.			-	order.		<b>⊠</b> e-mail
» Search O	ptions							
View Sessi	on History	Modify S	Search					
New Searc	<u>h</u>	( ( join <in< td=""><td>&gt;metadata ) <an< td=""><td>d&gt; ( merge<in< td=""><td>&gt;metadata ) )<a< td=""><td>and&gt; ( mir</td><td>ning<in>metadata)</in></td><td>Search</td></a<></td></in<></td></an<></td></in<>	>metadata ) <an< td=""><td>d&gt; ( merge<in< td=""><td>&gt;metadata ) )<a< td=""><td>and&gt; ( mir</td><td>ning<in>metadata)</in></td><td>Search</td></a<></td></in<></td></an<>	d> ( merge <in< td=""><td>&gt;metadata ) )<a< td=""><td>and&gt; ( mir</td><td>ning<in>metadata)</in></td><td>Search</td></a<></td></in<>	>metadata ) ) <a< td=""><td>and&gt; ( mir</td><td>ning<in>metadata)</in></td><td>Search</td></a<>	and> ( mir	ning <in>metadata)</in>	Search
		☐ Chec	k to search on	ly within this	results set	•		
» Key		Display	Format:	Citation C	Citation & A	Abstract		
IEEE JNL	IEEE Journal or Magazine	ر view so	elected items	Select Al	I Deselect Al	11		
IEE JNL	IEE Journal or Magazine	<b>V</b>		,		-		
IEEE CNF	IEEE Conference Proceeding	□ 1.			ssociation rul	les in re	lational database	es
IEE CNF	IEE Conference Proceeding		Houtsma, M.; Data Engineer		roceedings of t	the Eleve	enth International	Conference
IEEE STD	IEEE Standard		6-10 March 19 Digital Object		25 - 33 1109/ICDE.199	95 38041	13	
				Full Text: PD	<u>DF(</u> 608 KB) I			
		<b>2.</b>	microbial info Dawyndt, P.; \ Knowledge an Volume 17, Is	ormation sou Vancanneyt, I ad Data Engin ssue 8, Aug. Identifier 10.1 Full Text: PD	urces M.; De Meyer, neering, IEEE 2005 Page(s): 1109/TKDE.20	H.; Swin <u>Transact</u> :1111 - 1 005.131	ions on 126	es during t
		☐ 3.	Baoliang Liu; Computer and 21-23 Sept. 20	Jianzhong Li; I Information 005 Page(s):5 Identifier 10.1 Full Text: <u>PD</u>	Lei Nie; Yanq Technology, 2 58 - 64 1109/CIT.2005	iu Zhang <u>1005. CIT</u> 5.145	2005. The Fifth I	
		<b>4.</b>	Efficient reco Liang Jin; Che Database Sys International C 26-28 March 2 Digital Object AbstractPlus   Rights and Pe	en Li; Mehrotra tems for Adva Conference or 2003 Page(s): Identifier 10.1 Full Text: PD	ra, S.; <u>anced Applicat</u> <u>n</u> :137 - 146 1109/DASFAA	tions, 200		)3). Proceed

**Proceedings** 

Toyama, M.; Ohara, A.;

5. Hash-based symmetric data structure and join algorithm for OLAP applic

Database Engineering and Applications, 1999. IDEAS '99. International Sympo

2-4 Aug. 1999 Page(s):231 - 238 Digital Object Identifier 10.1109/IDEAS.1999.787272 AbstractPlus | Full Text: PDF(508 KB) IEEE CNF Rights and Permissions

Indexed by #Inspec\* Help Contact Us Privacy &: © Copyright 2006 IEEE -



Home | Login | Logout | Access Information | Alerts |

#### **Welcome United States Patent and Trademark Office**

□ Search Results	BROWSE	SEARCH	:	IEEE XPLORE GUIDE
Results for "( ( multiple <in>metadata ) <and> ( merg</and></in>	je <in>metadata ) )<ai< td=""><td>nd&gt; ( mining<in< td=""><td>**</td><td></td></in<></td></ai<></in>	nd> ( mining <in< td=""><td>**</td><td></td></in<>	**	

Search Res	sults		BROWSE	: \$	SEARCH	: 18	EE XPLORE GI	JIDE
Your search	"( ( multiple <in>metadata ) n matched 11 of 1320520 do n of 100 results are displayed</in>	cuments.			-			<b>⊠</b> e-mail
» Search O	ptions					•		
View Session	on History	Modify Sea						<b>-</b>
New Search	<u>h</u>	( ( multiple <i< td=""><td>n&gt;metadata ) <an< td=""><td>d&gt; ( merge&lt;</td><td>in&gt;metadata )</td><td>)<and> (</and></td><td>mining<in>metad</in></td><td>a Search</td></an<></td></i<>	n>metadata ) <an< td=""><td>d&gt; ( merge&lt;</td><td>in&gt;metadata )</td><td>)<and> (</and></td><td>mining<in>metad</in></td><td>a Search</td></an<>	d> ( merge<	in>metadata )	) <and> (</and>	mining <in>metad</in>	a Search
		Check	to search only w	ithin this res	sults set	:		
» Key		Display Fo	rmat:	ation C	Citation & Al	bstract		
IEEE JNL	IEEE Journal or Magazine	ட view sele	cted items   s	Select All 1	Deselect All			
IEE JNL	IEE Journal or Magazine	<b>*</b>						
IEEE CNF	IEEE Conference Proceeding	13	pproximate ont	,		semanti	ic Web	
IEE CNF	IEE Conference Proceeding	<u>Fı</u>	chardson, B.; Mazzy Information olume 2, 27-30 o	, 2004. Pro	cessing NAF		IEEE Annual Me	eeting of the
IEEE STD	IEEE Standard		gital Object Iden					
			ostractPlus   Fullights and Permis		(514 KB) IE	EEE CNI	F	
		m Di <u>Ki</u> Va Di	nowledge accur icrobial informa awyndt, P.; Vand nowledge and Di blume 17, Issue gital Object Iden ostractPlus   Full ghts and Permis	ation source canneyt, M., ata Enginee 8, Aug. 20 atifier 10.110 Text: <u>PDF</u> (	ces ; De Meyer, l ering, IEEE T 005 Page(s): 09/TKDE.200	H.; Swing <u>Fransacti</u> 1111 - 11 05.131	ons on 126	ies during t
		Sa Di or 26 Di At	arda, N.L.; Sriniv atabase and Exp	ras, N.V.; pert System Page(s):240 tiffer 10.110 Text: <u>PDF</u> (	s Application ) - 245 09/DEXA.199	ns, 1998. 98.70740	of association r	
		Yu Sc Ca 2- Di	educing uncerta the Li; Haihong I oftware Engineer onference 1997. 5 Dec. 1997 Pag gital Object Iden ostractPlus   Full ghts and Permis	Dai; ring Confere APSEC '97 ge(s):97 - 10 tifier 10.110 Text: <u>PDF</u> (	ence, 1997. / 7 and ICSC '9 05 09/APSEC.1	97. Proce	166	ational Com
		Ho	outsma, M.; Swa	mi, A.;			ational databas	

AbstractPlus | Full Text: PDF(608 KB) | IEEE CNF Rights and Permissions 6. Enhancing the semantic interoperability of multimedia through a core on Hunter, J.; Circuits and Systems for Video Technology, IEEE Transactions on Volume 13, Issue 1, Jan. 2003 Page(s):49 - 58 Digital Object Identifier 10.1109/TCSVT.2002.808088 AbstractPlus | References | Full Text: PDF(809 KB) | IEEE JNL Rights and Permissions 7. Merging Multiple Data Streams on Common Keys over High Performance Mazzucco, M.; Ananthanarayan, A.; Grossman, R.L.; Levera, J.; Bhagavantha Supercomputing, ACM/IEEE 2002 Conference 16-22 Nov. 2002 Page(s):67 - 67 Digital Object Identifier 10.1109/SC.2002.10044 AbstractPlus | Full Text: PDF(208 KB) | IEEE CNF Rights and Permissions 8. Knowledge Creation Opportunities in the Data Mining Process Brohman, M.K.; System Sciences, 2006. HICSS '06. Proceedings of the 39th Annual Hawaii Int Conference on Volume 8, 04-07 Jan. 2006 Page(s):170c - 170c Digital Object Identifier 10.1109/HICSS.2006.238 AbstractPlus | Full Text: PDF(176 KB) | IEEE CNF Rights and Permissions 9. Architecture of sensor based agricultural information system for effective П farm activities Chaudhary, S.; Sorathia, V.; Laliwala, Z.; Services Computing, 2004. (SCC 2004). Proceedings. 2004 IEEE Internationa 15-18 Sept. 2004 Page(s):93 - 100 Digital Object Identifier 10.1109/SCC.2004.1357994 AbstractPlus | Full Text: PDF(377 KB) IEEE CNF Rights and Permissions 10. Structured Knowledge Source Integration and its applications to informa Masters, J.; Information Fusion, 2002. Proceedings of the Fifth International Conference or Volume 2, 8-11 July 2002 Page(s):1340 - 1346 vol.2 Digital Object Identifier 10.1109/ICIF.2002.1020968 AbstractPlus | Full Text: PDF(510 KB) IEEE CNF Rights and Permissions 11. Association Pattern Language Modeling Chien J.-T.: IEEE Transactions on Audio, Speech, and Language Processing: Accepted for publication Volume PP, Issue 99, 2005 Page(s):1 - 10 Digital Object Identifier 10.1109/TSA.2005.858551

6-10 March 1995 Page(s):25 - 33

Digital Object Identifier 10.1109/ICDE.1995.380413

Help Contact Us Privacy &:

AbstractPlus | Full Text: PDF(384 KB) | IEEE JNL

Indexed by

© Copyright 2006 IEEE -



Home | Login | Logout | Access Information | Alerts |

#### Welcome United States Patent and Trademark Office

□ Search Session History

**BROWSE** 

**SEARCH** 

**IEEE XPLORE GUIDE** 

Thu, 23 Feb 2006, 10:15:31 AM EST

Edit an existing query or compose a new query in the Search Query Display.

# Select a search number (#)

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

<del></del>	

#### **Recent Search Queries**

Search Query Display

- #1 ((mining<in>metadata) <and> (template<in>metadata))
  <and> (monotonicity<in>metadata)

  ((mining<in>metadata) <and> (template<in>metadata))

  ((mining<in>metadata) <and> (template<in>metadata))

  <and> (pattern<in>metadata)
- #3 (( mining<in>metadata ) <and> ( template<in>metadata ) ) <and> ( attributes<in>metadata )
- #4 ((datasets<in>metadata) <and>(template<in>metadata))
  <and>(attributes<in>metadata)
- #5 ((irrelevant<in>metadata) <and> (template<in>metadata)) <and> (attributes<in>metadata)
- #6 ((join<in>metadata) <and> (template<in>metadata))<and> (attributes<in>metadata)
- #7 ((join<in>metadata)<and> (merge<in>metadata))<and> (attributes<in>metadata)</a>
- #8 ((join<in>metadata) <and> (merge<in>metadata))<and> (mining<in>metadata)
- #9 (( multiple<in>metadata ) <and> ( merge<in>metadata ) )<and>
   ( mining<in>metadata )

્રાંગવા ાર્યવાજસ્તાના . . . .

Indexed by Inspec

Help Contact Us Privacy &:

© Copyright 2006 IEEE -

We are experiencing some difficulties with our database and, consequently, some document links may be incorrect. Please bear with us while we fix the problem.

44 citations found. Retrieving documents...

Houtsma M., Swami A., "Set-Oriented Mining of Association Rules", Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, USA, October 1993

CiteSeer Home/Search Document Not in Database Summary Related Articles Check

This paper is cited in the following contexts:

Knowledge Discovery From Distributed And Textual Data - Cho (1999) (1 citation) (Correct)

....the large itemsets with confidence greater than the predetermined threshold. The first stage, which usually involves counting on the original database, is the most timeconsuming process and thus attracts many researchers efforts in speeding it up. Popular algorithms such as the AIS [5] and SETM [67] generate candidate itemsets on the fly while reading the data. These algorithms have the disadvantage of generating and counting too many candidate itemsets that turn out to be small. Agrawal and Srikant [3] later on modified the candidate itemset generation so that only itemsets found large in ....

Houtsma M. and Swami A., "Set-oriented Mining of Association Rules", Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, 1993.

Mining Association Rules From Market Basket Data.. - Hilderman.. (1998) (Correct)

....Machine Learning, Itemsets, Association Rules. 1 Introduction The problem of mining association rules from market basket data has recently been an important research topic in the area of knowledge discovery from databases. It was originally introduced in [2] and studied extensively in [1, 5, 25, 26, 31, 19, 23, 29, 30, 3, 4, 33, 14]. The problem is typically examined in the context of discovering buying patterns from retail sales transactions. Although there are many similar data mining applications which can be modelled in this way, we again study the problem using the retail store example because of its intuitive nature ....

....of share for itemsets, and redefine the notions of frequent itemsets and confidence. We refer to this extended formalism as the shareconfidence framework for association rules and refer to the new itemset measures as simply share measures. In this framework, any of the algorithms presented in [2, 3, 16, 19, 22, 23, 29, 30, 31, 32, 33] can used to generate frequent itemsets using our new definition for frequent itemset. The definitions in this section have been implemented in a data mining system for analyzing market basket data. This system is an extension of DB Discover, a software tool for knowledge discovery form databases ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. In Proceedings of the 11th International Conference on Data Engineering (ICDE'95), pages 25-34, 1995.

An Algorithm for Mining Association Rules Using Perfect.. - Özel, Güvenir (Correct)

....is greater than or equal to the minimum confidence. However the first step of association rule mining, finding the frequent itemsets, is very resource consuming task and for that reason, it has been one of the most popular research field in data mining. Several algorithms, AIS [3] SETM [8], Apriori [4] Direct Hashing and Pruning [5, 9] Partition [10] Sampling [11] and some other parallel algorithms [12] have been developed. In this study, a fast algorithm based on Direct Hashing and Pruning (DHP) algorithm is proposed. The DHP algorithm is described in Section II, our algorithm ....

M. Houtsma and A. Swami, "Set-Oriented Mining of Association Rules", Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, (Oct. 1993).

Intension Mining: A New Paradigm in Knowledge Discovery - Gupta, Bhatnagar, Wasan.. (2000) (Correct)

....cation with no change in semantics. Data Mining algorithms operate on previously selected, cleaned and transformed data. The choice of the mining algorithm depends on the type of knowledge to be discovered. Intra record links, termed Association Rules can be discovered using algorithms given in [4, 5, 6, 29, 31, 35, 41] etc. Database segmentation can be performed using various clustering techniques [18, 28, 37, 38, 51] etc. Classi cation can be performed by inducing either a decision tree or decision rule [1, 3, 8, 21, 23] or by neural network techniques [32] Presentation of the Discovered Knowledge is the ....

M. Houstma and A. Swami. Set oriented Mining of Association Rules. In Proceedings of the Int'l Conf. on Data Engineering, pages 25-33, 1995. 64

Discovering Interesting Association Rules in Medical Data - Ordonez, Santana, de Braal (2000) (Correct)

....are used with basket data. Medical data sets are more complex and thus present many new challenges. This paper incorporates some ideas from our previous work to mine rules on segmented images [16] Most papers published in the database literature concentrate on optimizing the first phase [18, 7, 12, 13, 14, 19, 17] but a few look at the problem of also improving rule generation (2nd phase) 6, 18, 15] For instance, 14] proposes an algorithm to summarize associations when they are too many. 7] attacks the problem of inserting transactions on an already mined set and proposes an algorithm that ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. Technical Report RJ 9567, IBM, October 1993.

Fast Algorithms for Mining Association Rules - Agrawal, Srikant (1994) (848 citations) (Correct)

....example, D could be a data file, a relational table, or the result of a relational expression. An algorithm for finding all association rules, henceforth referred to as the AIS algorithm, was presented in [AIS93b] Another algorithm for this task, called the SETM algorithm, has been proposed in [HS93] In this paper, we present two new algorithms, Apriori and AprioriTid, that differ fundamentally from these algorithms. We present experimental results, using both synthetic and real life data, showing that the proposed algorithms always outperform the earlier algorithms. The performance gap is ....

....person who orders a comforter also orders a flat sheet, a fitted sheet, a pillow case, and a ruffle. The algorithms in Section 3 generate suchmulti consequent rules. In Section 4, we show the relative performance of the proposed Apriori and AprioriTid algorithms against the AIS [AIS93b] and SETM [HS93] algorithms. To make the paper self contained, we include an overview of the AIS and SETM algorithms in this section. We also describe how the Apriori and AprioriTid algorithms can be combined into a hybrid algorithm, AprioriHybrid, and demonstrate the scale up properties of this algorithm. We ....

[Article contains additional citation context not shown here]

Maurice Houtsma and Arun Swami. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden ResearchCenter, San Jose, California, October 1993. 30

SQL Based Association Rule Mining using Commercial.. - Yoshizawa.. (2000) (Correct)

....and credit card fraud indications are widely recognized. One method of data mining is finding association rule [1] Basket data analysis is typical of this method. There are some approaches proposed to mine association rules, 1,2,6,9] some of them are based on relational database standard SQL [3,7,8]. But this kind of mining is known as CPU power demanding application and it has to handle very large amounts of transaction data. Unfortunately SQL approach is reported to have drawback in performance although it has many advantages such as seamless integration with existing system and high ....

....required by association rule mining. This fact motivated us to examine how efficiently SQL based association rule mining can be parallelized and speeded up using commercial parallel database system (IBM DB2 UDB EEE) We propose two techniques to enhance association rule mining query based on SETM [3]. And we have also compared the performance with commercial mining tool (IBM Intelligent Miner) Our performance evaluation shows that we can achieve comparable performance with commercial mining tool using only 4 nodes. Some considerable works on effective SQL queries to mine association rule ....

[Article contains additional citation context not shown here]

M. Houtsma, A. Swami. Set-oriented Mining of Association Rules. In Proc. of International Conference on Data

Engineering (ICDE), 1995.

Extended Concepts for Association Rule Discovery - Rantzau (1997) (Correct)

....at most twice. In the first scan it generates all candidates and in the second their support is computed. Apriori outperforms Partition only when the minimum support threshold is set high. The Partition algorithm lends itself to an implementation on parallel computers. The SETM algorithm [HS95, AS94a] uses SQL to generate frequent itemsets. Like AIS, candidates are generated while transactions are read from the database. However, SETM separates candidate generation from counting. It has a worse performance than AIS for both synthetic and real life datasets. The algorithm presented in ....

Maurice Houtsma and Arun Swami. Set-oriented Mining of Association Rules. In Proceedings of the 11th International Conference on Data Engineering, Taipei, Taiwan, pages 25–33, March 1995.

<u>Performance Evaluation and Optimization of Join Queries.. - Thomas, Chakravarthy (1998) (2 citations) (Correct)</u>

....objectrelational extensions to execute mining operations. This entails transforming the mining operations into database queries and in some cases developing newer techniques that are more appropriate in the database context. The UDF based (user de ned function) approach in [2] the SETM algorithm [5], the formulation of association rule mining as query ocks [10] and SQL queries for mining [9] all belong to this category. Two categories of SQL implementations for association rule mining one based purely on SQL 92 and the other using the object relational extensions to SQL (SQL OR) are ....

....= q.tid and q.tid = r.tid We can also use the Subquery approach to generate T 3 if that is less expensive. T 3 will contain exactly the same tuples produced by subquery Q 3. The Set oriented Apriori algorithm bears some resemblance with the three way join approach in [9] the SETM algorithm in [5] and the AprioriTid algorithm in [3] In the three way join approach, the temporary table T k stores for each transaction, the identi ers of the candidates it supported. T k is generated by joining two copies of T k 1 with C k. The generation of F k requires a further join of T k with C k. The ....

[Article contains additional citation context not shown here]

M. Houtsma and A. Swami. Set-oriented mining of association rules. In Int'l Conference on Data Engineering, Taipei, Taiwan, March 1995.

Algorithms For Computing Association Rules Using A. - Graham Goulbourne Frans (2000) (1 citation) (Correct)

....support for all members of C k, and from this, produces the set L k of interesting sets of size k. This is then used to derive the candidate sets C k 1, using the downward closure property, that all the subsets of any member of C k 1 must be members of L k. Other algorithms, AIS [1] and SETM [3], have the same general form but differ in the way the candidate sets are derived. Two aspects of the performance of these algorithms are of concern: the number of passes of the database that are required, which will in general be one greater than the number of attributes in the largest ....

Houtsma, M. and Swami, A. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Centre, San Jose, October 1993.

<u>Performance Evaluation and Optimization of Join Queries.. - Thomas, Chakravarthy (1998) (2 citations) (Correct)</u>

....extensions to execute mining operations. This entails transforming the mining operations into database queries and in some cases developing newer techniques that are more appropriate in the database context. The UDF based (user de ned function) approach in [AS96] the SETM algorithm [HS95] the formulation of association rule mining as query ocks [TUA 98] and SQL queries for mining [STA98] all aim at tighter database integration. STA98] presents a detailed study of the various architectural alternatives for mining data stored in a DBMS. It has been reported that for ....

....Section 3.3. Figure 11: Comparison of Subquery and Set oriented Apriori approaches In Figure 11, we show

the relative performance of Subquery and Set oriented Apriori approaches for the two datasets. The chart shows the total time taken for each of the di erent passes. We ran the SETM algorithm [HS95] also for a few support values and found that it is an order of magnitude slower. Set oriented Apriori performs better than Subquery for all the support values. The rst two passes of both the approaches are similar and they take approximately equal amount of time. The di erence between ....

Maurice Houtsma and Arun Swami. Set-oriented mining of association rules. In Int'l Conference on Data Engineering, Taipei, Taiwan, March 1995.

Incremental Mining of Constrained Associations - Thomas, Chakravarthy (1998) (6 citations) (Correct)

....frequent itemsets) generated at each level is based on the observation that if an itemset S appears in c baskets, then any subset of S appears in at least c baskets. The need for applying association rule mining to data stored in databases data warehouses has motivated researchers to [SK97, HS95, AS96, STA98, TC98, TS98] i) study alternative architectures for mining over data stored in databases, ii) translate association rule mining algorithms to work with relational and object relational databases, ii) optimize the mining algorithms beyond what the current relational query optimizers ....

Maurice Houtsma and Arun Swami. Set-oriented mining of association rules. In Int'l Conference on Data Engineering, Taipei, Taiwan, March 1995.

Parallel Mining of Association Rules - Agrawal, Shafer (1996) (52 citations) (Correct)

....data mining is that it will deliver technology that will enable development of a new breed of decision support applications. Discovering association rules is an important data mining problem [1] Recently, there has been considerable research in designing fast algorithms for this task [1] 3] 5] [6] [8] 12] 9] 11] However, with the exception of [10] the work so far has been concentrated on designing serial algorithms. Since the databases to be mined are often very large (measured in gigabytes and even terabytes) parallel algorithms are required. We present in this paper three parallel ....

....upon the patterns the different transactions support. This algorithm also incorporates load balancing. These algorithms are based upon the serial algorithm Apriori which was first presented in [3] We chose the Apriori algorithm because of its superior performance over the earlier algorithms [1] [6], as shown in [3] We preferred Apriori over AprioriHybrid, a somewhat faster algorithm in [3] because AprioriHybrid is harder to parallelize; the performance of AprioriHybrid is sensitive to heuristically determined parameters. Furthermore, the performance of Apriori can be made to approximate ....

[Article contains additional citation context not shown here]

Maurice Houtsma and Arun Swami. Set-oriented mining of association rules. In Int'l Conference on Data Engineering, Taipei, Taiwan, March 1995.

Parallel SQL Based Association Rule Mining on.. - Pramudiono.. (1999) (Correct)

....called large itemsets. 2. **Generate** the desired rules using large itemsets. **Since** the first step consumes most of processing time, development of mining algorithm has been concentrated on this step. **In our experiment we employed ordinary standard SQL query that is similar to SETM algorithm [3].It is shown in figure 1. <b>CREATE** TABLE SALES (id int, item int) PASS 1 CREATE TABLE C 1 (item 1 int, cnt int) CREATE TABLE R 1 (id int, item 1 int) INSERT INTO C 1 SELECT item AS item 1, COUNT( FROM SALES GROUP BY item HAVING COUNT( min support; INSERT INTO R 1 SELECT ....

M. Houtsma, A. Swami. Set- oriented Mining of Association Rules. In Proc. of International Conference on Data Engineering, 1995.

Pincer-Search: An Efficient Algorithm for Discovering the.. - Lin, Kedem (1999) (Correct)

....Frequent Set Discovery We briefly discuss existing frequent set discovery algorithms in a roughly chronological order. AIS and SETM Algorithms The problem of association rule mining was first introduced in [2] An algorithm called AIS was given for discovering the frequent set. SETM algorithm [13] was later designed to use only

standard SQL commands to find the frequent set. The Apriori algorithm [3] described above, performs much better than AIS and SETM. The OCD Algorithm It is worth adding, that concurrently with the Apriori algorithm, OCD algorithm [19] used the same closure property ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Center, Oct. 1993.

Computing Association Rules Using Partial Totals - Graham Goulbourne Frans (2001) (2 citations) (Correct)

....support for all members of C k, and from this, produces the set L k of interesting sets of size k. This is then used to derive the candidate sets C k 1, using the downward closure property, that all the subsets of any member of C k 1 must be members of L k. Other algorithms, AIS [1] and SETM [3], have the same general form but differ in the way the candidate sets are derived. Two aspects of the performance of these algorithms are of concern: the number of passes of the database that are required, which will in general be one greater than the number of attributes in the largest ....

Houtsma, M. and Swami, A. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Centre, San Jose, October 1993. 10

Fast Algorithms for Discovering the Maximum Frequent Set - Lin (1998) (Correct)

....Itemsets f1,2,3,5g and f1,2,5g were not considered, since the item 5 was not in the transaction. **Two** complicated heuristics, remaining tuples optimization and pruning function optimization, were used to prune candidates. **Unfortunately**, this algorithm still generates too many candidates. **SETM [HS93] algorithm was later designed to use only standard SQL commands to find the frequent set.** However, like AIS, SETM also creates candidates on the fly while reading the database. **Both** algorithms are not efficient, since they generate and count too many unnecessary candidates. 2.4.2 Apriori and OCD ....

M. Houtsma and A. Swami. *Set-oriented mining of association rules*. Research Report RJ 9567, IBM Almaden Research Center, Oct. 1993.

Integrating Data Mining with Relational DBMS: A. - Nestorov, Tsur (1999) (1 citation) (Correct)

.... that requires consideration: can we achieve a comparable, or at least an acceptable level of performance from these integrated methods when compared to the special purpose external methods This question was previously examined in a more narrow context of association rules and a particular DBMS in [7] and [2] Section 2 of this paper elaborates on the general architectural choices available and their comparison. The idea of flocks [11] was presented as a framework for performing complex data analysis tasks on relational database systems. The method consists of a generator of candidate query ....

H. Houtsma and A. Swami. Set-oriented mining of association rules. In Proceedings of International Conference on Data Engineering, pages 25-33, Taipei, Taiwan, March 1995.

Efficient Mining for Association Rules with Relational.. - Rajamani, Cox, Iyer, al. (1999) (3 citations) (Correct)

....With the (Transaction id, Item) schema the Transaction id value would be repeated for every item bought in that transaction. The SC data model would be useful for performing conventional relational queries against items bought in transactions. Some of the early work in association rule mining [11] propose the use of such relational queries for discovering association rules, and work with this data model. However, later work [3] have shown significant performance improvement by using Apriori based algorithms that did not use relational queries in their implementation. But, to the best of ....

....proposed for inter operability in multi database systems and not for providing the flexibility and functionality required by data mining applications. **Agrawal** and Shim [2] show the benefit of using UDFs for the development of applications tightly coupled with the database engine. **Houtsma and Swami [11] proposed SETM, an SQL based algorithm for association rule mining.** Their algorithm uses simple database operations sorting and mergescan joins. **However**, their joins are more expensive as they are against the input data table and they do not have an efficient candidate set pruning such as ....

M. Houtsma and A. Swami. Set-oriented Mining of Association Rules. Technical Report RJ 9567, IBM Almaden Research Center, October 1993.

Mining Association Rules - Cengiz (Correct)

....with predefined template patterns. If a found rule matches to any pattern than it is said to be interesting, otherwise uninteresting. 4 Serial Algorithms Since [AIS93] the introduction of the problem, there has been work on designing algorithms for mining Boolean association rules [AS94] [HS95] [MTV94] SON95] PCY95] The algorithm presented in [AIS93] incorporates buffer management, novel estimation, and pruning techniques to provide efficiency. This algorithm finds only Boolean association rules. This work was later extended to cover quantitative rules [SA95] HF95] Those works try ....

Maurice Houtsma and Arun Swami. Set-oriented mining of association rules. In Int'l Conference on Data Engineering, Taipei, Taiwan, March 1995.

Mining for Complex Patterns in Large Databases Using.. - Jermaine, Miller (1998) (Correct)

....the model representations upon which they rely. The limitations include restrictions on the type and characteristics of the data, limiting the accuracy and completeness of the models. 3. 1 Association Rules Association rules have been the most prevalent mining method in the database literature [1, 2, 3, 4, 9, 11, 12, 14, 15, 20, 22, 23]. While very useful in describing certain large scale trends or relationships among the data in very large datasets, association rules are unable to describe certain relationships that have very strong statistical significance but occur in a only a relatively small subset of the data precisely ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. In Proceedings of the International Conference on Data Engineering, pages 25-34, 1995.

Parallel Data Mining for Association Rules on.. - Zaki, Ogihara.. (1996) (30 citations) (Correct)

....large itemsets, and in the second pass the support for all these is measured. The above algorithms are all specialized black box techniques which do not use any database operations. Algorithms using only general purpose DBMS systems and relational algebra operations have also been proposed [7, 8]. The work closest to this from the machine learning literature is the KID3 algorithm presented in [12] The main problem with their approach is that it may take exponential time in the worst case as opposed to the polynomial time algorithms presented in the above papers. There has been very ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. In RJ 9567. IBM Almaden, Oct. 1993.

Formal Logics of Discovery and Hypothesis Formation By Machine - Hájek, Holena (Correct)

....tables [5, 9, 31, 32, 55, 56] Scope. GUHA relates, in particular, to mining association rules. Indeed, if A = fA 1; Am g is the set of binary attributes in a database of size k, and if X;Y ae A; X·Y = then the association rule X) Y is significant in the database (according to [1, 2, 27, 28, 40, 47, 54]) if and only if the GUHA sentence i2X A i =B;p i2Y A i holds for the k Theta m dichotomous data matrix formed by the values of the attributes from A. Here, B;p is a founded version of the generalized quantifier = p mentioned in section 2 (version requiring the frequence a to be at ....

Houtsma, M., and Swami, A. Set-oriented mining of association rules. Tech. rep., IBM Almaden Research Center, 1993.

Beyond Market Baskets: Generalizing Association Rules.. - Silverstein, Brin.. (1997) (40 citations) (Correct)

....properties, text mining 1. Introduction One particularly well studied problem in data mining is the search for association rules in market basket data (Agrawal et al. 1993a, Agrawal et al. 1993b, Klemettinen et al. 1994, Mannila et al. 1994, Agrawal and Srikant, 1994, Han and Fu, 1995, Houtsma and Swami, 1995, Park et al. 1995, Srikant and Agrawal, 1995, Savasere et al. 1995, Agrawal et al. 1996, Toivonen, 1996) In this setting, the base information consists of register transactions of retail stores. The goal is to discover buying patterns such as two or more items that are bought together often. ....

Houtsma, M. and Swami, A. 1995. Set-oriented mining of association rules. Proceedings of the International Conference on Data Engineering, pp. 25–34.

Discovering all Most Specific Sentences by Randomized.. - Gunopulos, al. (1997) (31 citations) (Correct)

....these candidates, the algorithm calls the function q to check whether really belongs to Th. **This** iterative procedure is performed until no more sentences in Th are found. **This** level wise algorithm has been used in various forms in finding association rules, episodes, sequential rules, etc. **[2, 3, 24, 23, 1, 13, 14, 25].** The drawback with this algorithm is that it always computes the whole set Th(L; r; q) even in the cases where a condensed representation of Th using most specific sentences would be useful. **Given** Th, a sentence 2 Th is a most specific sentence of Th, if for no 2 Th we have OE. Denote ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, October 1993.

Query Flocks: A Generalization of Association-Rule Mining - Tsur, Ullman.. (1998) (43 citations) (Correct)

....answer to a question like find all the pairs of items that appear in at least c market baskets. 1.3 The Problem With SQL as a Mining Language In principle, we can express a query about pairs of items that appear in a large number of baskets in conventional SQL. **This approach was examined by [HS95], for instance.** The problem is that the right optimizations are beyond the state of the art in commercial database systems. **For** example, Fig. 1 shows how to express the query find all pairs of items that appear together in at least 20 market baskets. 1 There, we join baskets with itself, ....

M. Houtsma and A. Swami, "Set-oriented mining of association rules," Proc. Intl. Conf. on Data Engineering, pp. 25--34.

Query Flocks: A Generalization of Association-Rule Mining - Dick Tsur (1998) (43 citations) (Correct)

....answer to a question like find all the pairs of items that appear in at least c market baskets. 1.3 The Problem With SQL as a Mining Language In principle, we can express a query about pairs of items that appear in a large number of baskets in conventional SQL. **This approach was examined by [HS95], for instance.** The problem is that the right optimizations are beyond the state of the art in commercial database systems. **For** example, Fig. 1 shows how to express the query find all pairs of items that appear together in at least 20 market baskets. 1 There, we join baskets with itself, ....

M. Houtsma and A. Swami, "Set-oriented mining of association rules," Proc. Intl. Conf. on Data Engineering, pp. 25-34.

Mining for Strong Negative Associations in a Large Database.. - Ashok Savasere (1998) (6 citations) (Correct)

.... of the promising areas of research encompassing databases, statistics and machine learning [6, 12, 15] Recently, there has been considerable interest in finding associations between items in a database of customer transactions, such as the sales data collected at super market check out counters [1, 2, 5, 7, 11, 8, 14, 4]. Association rules identify items that are most often bought along with certain other items by a significant fraction of the customers. For example, we may find that 95 of the customers who bought bread also bought milk. Every rule must satisfy two user specified constraints: one is a measure ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. In Proceedings of the International Conference on Data Engineering, Taipei, Taiwan, March 1995.

Mining Quantitative Association Rules in Large Relational Tables - Srikant, Agrawal (1996) (201 citations) (Correct)

....supports are counted are new. We present our experience with this solution on a real life dataset in Section 6. 1. 3 Related Work Since the introduction of the (Boolean) Association Rules problem in [AIS93] there has been considerable work on designing algorithms for mining such rules [AS94] [HS95] [MTV94] SON95] PCY95] This work was subsequently extended to finding association rules when there is a taxonomy on the items in [SA95] HF95] Related work also includes [PS91] where quantitative rules of the from x = q x)

y = q y are discovered. However, the antecedent and consequent are ....

Maurice Houtsma and Arun Swami. Set-oriented mining of association rules. In Int'l Conference on Data Engineering, Taipei, Taiwan, March 1995.

Fast Sequential and Parallel Algorithms for Association Rule. - Mueller (1995) (44 citations) (Correct)

....paper, rules were restricted to one item in the consequent but allowed any union of items in the antecedent.

This limitation was not justified by the algorithm itself because it finds all frequent sets, and this information is enough to produce rules without this limitation. 2.3. 2 SETM SETM [22] is designed to use only standard database operations to find frequent sets. For this reason, it uses its own data representation that stores every itemset supported by a transaction along with the transaction s TID. Figure 2.3 shows part of an example run of SETM on a tiny database and ....

Houtsma and Arun Swami. Set-oriented mining of association rules. Technical Report RJ 9567, IBM Research Report, Oct. 1993.

An Effective Hash-Based Algorithm for Mining Association Rules - Park, Yu (1995) (114 citations) (Correct)

....are referred to as large itemsets. **Once** all large itemsets are obtained, the desired association rules can be generated in a straightforward manner. **Subsequent** work in the literature followed this approach and focused on the large itemset generations. **Various algorithms have been proposed [3, 5, 8] to discover the large itemsets.** Generally speaking, these algorithms first construct a candidate set of large itemsets based on some heuristics, and then discover the subset that indeed contains large itemsets. **This** process can be done iteratively in the sense that the large itemsets discovered ....

M. Houtsma and A. Swami. Set-Oriented Mining of Association Rules. Technical Report RJ 9567, IBM Almaden Research Laboratory, San Jose, CA, October 1993.

Integrating Association Rule Mining with Relational Database.. - Sarawagi (1998) (46 citations) (Correct)

....proposal makes use of user defined functions (UDFs) in SQL statements to selectively push parts of the computation into the database system. The objective was to avoid one at a time record retrieval from the database, saving both the copying and process context switching costs. The SETM algorithm [10] for finding association rules was expressed in the form of SQL queries. However, as shown in [3] SETM is not efficient and there are no results reported on running it against a relational DBMS. Recently, the problem of expressing the association rules algorithm in SQL has been explored in [20] ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. In Int'l Conference on Data Engineering, Taipei, Taiwan, March 1995.

Evaluation of Sampling for Data Mining of Association Rules - Zaki, Parthasarathy, Li.. (1996) (16 citations) (Correct)

....large itemsets, and in the second pass the support for all these is measured. The above algorithms are all specialized black box techniques which do not use any database operations. Algorithms using only generalpurpose DBMS systems and relational algebra operations have also been proposed [5, 6]. In this paper we will use the Apriori algorithm to evaluate the effectiveness of sampling for data mining. We would like to point out that our results are about sampling, and as such independent of, and equally applicable to any algorithm for mining of association rules. We now present a brief ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. In RJ 9567. IBM Almaden, Oct. 1993.

Fast Algorithms for Mining Association Rules - Agrawal, Srikant (1994) (848 citations) (Correct)

....example, D could be a data file, a relational table, or the result of a relational expression. An algorithm for finding all association rules, henceforth referred to as the AIS algorithm, was presented in [AIS93b] Another algorithm for this task, called the SETM algorithm, has been proposed in [HS93] In this paper, we present two new algorithms, Apriori and AprioriTid, that differ fundamentally from these algorithms. We

present experimental results, using both synthetic and real life data, showing that the proposed algorithms always outperform the earlier algorithms. **The** performance gap is ....

....person who orders a comforter also orders a flat sheet, a fitted sheet, a pillow case, and a ruffle. The algorithms in Section 3 generate such multi consequent rules. In Section 4, we show the relative performance of the proposed Apriori and AprioriTid algorithms against the AIS [AIS93b] and SETM [HS93] algorithms. To make the paper self contained, we include an overview of the AIS and SETM algorithms in this section. We also describe how the Apriori and AprioriTid algorithms can be combined into a hybrid algorithm, AprioriHybrid, and demonstrate the scale up properties of this algorithm. We ....

[Article contains additional citation context not shown here]

Maurice Houtsma and Arun Swami. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, October 1993.

An Efficient Algorithm for Mining Association Rules in Large.. - Ashok Savasere (1995) (16 citations) (Correct)

....rule may state that 95 of customers who bought items A and B also bought C and D. **This** type of information may be used to decide catalog design, store layout, product placement, target marketing, etc. **Many algorithms have been discussed in the literature for discovering association rules [1, 8, 2].** One of the key features of all the previous algorithms is that they require multiple passes over the database. **For** disk resident databases, this requires reading the database completely for each pass resulting in a large number of disk reads. **In** these algorithms, the effort spent in performing ....

....is stopped when in some iteration n, no large itemsets are generated. The algorithm, in this case, makes n database scans. 2.1 Previous Work The problem of generating association rules was first introduced in [1] and an algorithm called AIS was proposed for mining all association rules. In [8], an algorithm called SETM was proposed to solve this problem using relational operations in a relational database environment. In [2] two new algorithms called Apriori and AprioriTid were proposed. These algorithms achieved significant improvements over the previous algorithms and were ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. Technical Report RJ 9567, IBM, October 1993.

Fast Algorithms for Mining Association Rules - Agrawal, Srikant (1994) (848 citations) (Correct)

....D. For example, D could be a data file, a relational table, or the result of a relational expression. An algorithm for finding all association rules, henceforth referred to as the AIS algorithm, was presented in [4] Another algorithm for this task, called the SETM algorithm, has been proposed in [13]. In this paper, we present two new algorithms, Apriori and AprioriTid, that differ fundamentally from these algorithms. We present experimental results showing that the proposed algorithms always outperform the earlier algorithms. The performance gap is shown to increase with problem size, and ....

....subsets of I to generate rules with multiple consequents. **Due** to lack of space, we do not discuss this subproblem further, but refer the reader to [5] for a fast algorithm. **In Section 3, we show the relative performance of the proposed Apriori and AprioriTid algorithms against the AIS [4] and SETM [13] algorithms.** To make the paper self contained, we include an overview of the AIS and SETM algorithms in this section. **We** also describe how the Apriori and AprioriTid algorithms can be combined into a hybrid algorithm, AprioriHybrid, and demonstrate the scaleup properties of this algorithm. **We** ....

[Article contains additional citation context not shown here]

M. Houtsma and A. Swami. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, October 1993.

<u>Levelwise Search and Borders of Theories in Knowledge Discovery - Mannila, Toivonen (1997)</u> (85 citations) (Correct)

....the rule is fr(X [fAg) fr(X)) All rules with frequency higher than a given threshold can be found eoeectively by using a simple algorithm for nding all frequent sets. A set X R is frequent, if fr(X) exceeds the given threshold.

Several algorithms for nding frequent sets have been presented [1, 2, 11, 14, 15, 16, 31, 35, 36, 37, 38]. The problem of nding all frequent sets can be described in our framework as follows. The description language L consists of all subsets X of elements of R. The selection predicate q(r; X) is true if and only if fr(X) minfr, where minfr is the frequency threshold given by the user. In the ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, October 1993.

Parallel Data Mining for Association Rules on.. - Zaki, Ogihara.. (1996) (30 citations) (Correct)

....large itemsets, and in the second pass the support for all these is measured. The above algorithms are all specialized black box techniques which do not use any database operations. Algorithms using only general purpose DBMS systems and relational algebra operations have also been proposed [5, 6]. There has been very limited work in parallel implementations of association algorithms. In [9] a parallel implementation of the DHP algorithm [8] is presented. However only simulation results on a shared nothing or distributed memory machine like IBM SP2 were presented. Parallel implementations ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. In RJ 9567. IBM Almaden, Oct. 1993.

Discovering Association Rules based on Image Content - Carlos Ordonez (1999) (8 citations) (Correct)

....content by Blobworld. 3.1. **Data** mining based on association rules At this point, we will consider in detail, the problem of finding associations. **The** problem of generating association rules was first introduced in [1] and an algorithm called AIS was proposed for mining all association rules. **In [13], an algorithm called SETM was proposed to solve this problem using relational operations.** In [2] two algorithms called Apriori and Apriori Tid were proposed. **These** algorithms achieved significant improvements over the previous algorithms. **The** rule generation process was also extended to include ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. Technical Report RJ 9567, IBM, October 1993.

Mining Generalized Association Rules - Srikant, Agrawal (1995) (255 citations) (Correct)

....) Hiking Boots and Clothes) Hiking Boots may Clothes Jackets Footwear Ski Pants Shirts Outerwear Shoes Hiking Boots Figure 1: Example of a Taxonomy not. The former may not have minimum support, and the latter may not have minimum confidence. Earlier work on association rules [1] 2] [5] [6] 7] did not consider the presence of taxonomies and restricted the items in association rules to the leaf level items in the taxonomy. However, finding rules across different levels of the taxonomy is valuable since: ffl Rules at lower levels may not have minimum support. Few people may buy ....

....T 0, where T 0 contains all the items in T as well as all the ancestors of each items in T. For example, if the transaction contained Jackets, we would add Outerwear and Clothes to get the extended transaction. We can then run any of the algorithms for mining association rules [1] 2] [5] [6] 7] on the extended transactions to get generalized association rules. However, this Basic algorithm is not very fast; two more sophisticated algorithms that we propose run 2 to 5 times faster than Basic (and more than 100 times faster on one real life dataset) We describe the Basic ....

[Article contains additional citation context not shown here]

M. Houtsma and A. Swami. Set-oriented mining of association rules. In Int'l Conference on Data Engineering, Taipei, Taiwan, March 1995.

<u>Evaluation of Sampling for Data Mining of Association Rules - Zaki, Parthasarathy, Li.. (1996)</u> (16 citations) (Correct)

....a certain user specified frequency, called minimum support. The second step consists of forming implication rules among the large itemsets [3] In this paper we only deal with the computationally intensive first step. Many algorithms for finding large itemsets have been proposed in the literature [1, 7, 3, 10, 12, 6, 13, 2]. In this paper we will use the Apriori algorithm [2] to evaluate the effectiveness of sampling for data mining. We chose

Apriori since it fast and has excellent scale up properties. **We** would like to observe that our results are about sampling, and as such independent of the mining algorithm ....

....algorithm [13] minimizes I O by scanning the database only twice. In the first pass it generates the set of all potentially large itemsets, and in the second pass their support is obtained. Algorithms using only general purpose DBMS systems and relational algebra operations have also been proposed [6, 7]. A theoretical analysis of sampling (using Chernoff bounds) for association rules was presented in [2, 10] We look at this problem in more detail empirically, and compare theory and experimentation. In [8] the authors compare sample selection schemes for data mining. They make a claim for ....

M. Houtsma and A. Swami. Set-oriented mining of association rules. In RJ 9567. IBM Almaden, Oct. 1993.

Itemset Materializing for Fast Mining - Of Association Rules (Correct)

No context found.

Houtsma M., Swami A., "Set-Oriented Mining of Association Rules", Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, USA, October 1993

Similarity Testing Between Heterogeneous Basket Datasets - Li, al. (2002) (Correct)

No context found.

M. Houtsma and A. Swami. *Set-oriented mining of association rules*. Research Report RJ 9567, IBM Almaden Research Center, San Jose, CA, October 1993.

MINTO: A Software Tool for Mining Manufacturing Databases - Haritsa (Correct)

No context found.

M. Houtsma and A. Swami, "Set-Oriented Mining Association Rules", International Conference on Data Engineering, March 1995.

CiteSeer.IST - Copyright Penn State and NEC

Sign in



Web Images Groups News Froogle Local more »

discovery relational data mining query templat Search Preferences

# Web Results 1 - 10 of about 19,400 for discovery relational data mining query template pruning procedure

### [PDF] Contents

File Format: PDF/Adobe Acrobat - <u>View as HTML</u>
We now focus on specific **data mining** and **pattern discovery** tasks that are often ... straints on **pattern templates**, such as wildcards with user-specified ... maya.cs.depaul.edu/~mobasher/papers/IC-Handbook-04.pdf - <u>Similar pages</u>

### [PDF] Untangling Text Data Mining

File Format: PDF/Adobe Acrobat - View as HTML of discovery of trends and patterns within tex-. tual data for more general purpose usage. ... port for issuing sequences of queries and related ... www.sims.berkeley.edu/~hearst/papers/acl99.pdf - Similar pages

## **Untangling Text Data Mining**

Instead, data mining applications tend to be (semi)automated discovery of trends ... sequences of queries and related operations across text collections, ... www.sims.berkeley.edu/ ~hearst/papers/acl99/acl99-tdm.html - 40k - Cached - Similar pages

### [PDF] Querying Multiple Sets of Discovered Rules

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> query languages. One of the earliest data mining query language. is the one based on templates [15]. In this technique, the user uses. a template to specify ... www.cs.uic.edu/~liub/publications/kdd2002-RuleQL.pdf - <u>Similar pages</u>

# [РDF] An Indexing Scheme for Fast Similarity Search in Large Time Series ... File Format: PDF/Adobe Acrobat - View as HTML

We begin by placing a **template** grid over the **query** Q. This **template** is exactly the same as the one ... Knowledge **Discovery** and **Data Mining**. pp 16-22, AAAI ... www.ics.uci.edu/~pazzani/Publications/ssdb99.pdf - <u>Similar pages</u>

#### [PDF] APA 5th Edition Template

File Format: PDF/Adobe Acrobat - View as HTML
The research on preference based frequent page

The research on preference-based frequent **pattern mining** is **related** to two categories of. previous studies: preference **queries** and frequent **pattern mining** ... www.cs.sfu.ca/~jpei/publications/prefmine-jdwm.pdf - <u>Similar pages</u>

#### [PDF] Final Report

File Format: PDF/Adobe Acrobat - View as HTML
David I Hand: Pattern detection and discovery

David J. Hand: **Pattern** detection and **discovery**. **Data mining** comprises ... background knowledge and us a **template**-based **mining** language to exploit it during ... www.esf.org/generic/1362/0110PReport.pdf - <u>Similar pages</u>

# Mining Software Engineering Data: A Survey - Tool Descriptions

AIRA includes different templates that can be used, one based on frames, ... Provides Discovery of Data Patterns - Yes Data Mining in CrossTables Analysis ... www.dacs.dtic.mil/techs/datamining/descriptions.shtml - 184k - Cached - Similar pages

# Citations: Finding Interesting Rules from Large Sets of Discovered ...

... of data mining to discover interesting relationships among attributes in ... Templates [10], pattern expressions that describe the form of rules that ...

discovery relational data mining query template pruning procedure pattern template - Goo... Page 2 of 2

citeseer.ist.psu.edu/context/35674/0 - 35k - Cached - Similar pages

[PDF] Inspec Alerts

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> sizeable set of **pattern templates** to be generated for **mining** or. **query**. ... Integrating **query** processing and **data mining** in. **relational** DBMSs ... inspec.iee.org/inspec/currentawareness/ pdfs/products/samples/IEE\_data\_mining.pdf - <u>Similar pages</u>

Try your search again on Google Book Search

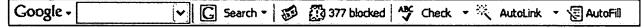
Gooooooogle >

Result Page:

1 2 3 4 5 6 7 8 9 10

<u>Next</u>

Free! Get the Google Toolbar. Download Now - About Toolbar



discovery relational data mining que

Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

Sign in



Web <u>Images</u> Groups News Froogle Local more »

anti monotonicity property discovery relational Search

Web Results 1 - 10 of about 479 for anti monotonicity property discovery relational data mining query tem

# (PDF) APA 5th Edition Template

File Format: PDF/Adobe Acrobat - View as HTML

Input: a mining query with a preference, the preference is either monotonic or anti-

monotonic; Output: the complete set of preference patterns; Method: ...

www.cs.sfu.ca/~ipei/publications/prefmine-jdwm.pdf - Similar pages

# [PDF] Pushing Convertible Constraints in Frequent Itemset Mining

File Format: PDF/Adobe Acrobat - View as HTML

convertible anti-monotone wrt value descending order (see also Example 3). This

property. provides great flexibility in data mining query optimization. ...

www.cs.sfu.ca/~jpei/publications/conv jn.pdf - Similar pages

# [PDF] Simultaneous Optimization of Complex Mining Tasks with a ...

File Format: PDF/Adobe Acrobat - View as HTML

vision that a Knowledge Discovery and Data Mining System (KD- ... issuing a sequence of

related complex mining queries, and 2) Multi-...

www.cs.kent.edu/~jin/Papers/mqo.pdf - Similar pages

# [PDF] InfoMiner: Mining Surprising Periodic Patterns

File Format: PDF/Adobe Acrobat - View as HTML

Periodic pattern discovery is an important problem in mining time series data and has

wide ... (1) The surprise of a pattern occurrence is anti-monotonic ...

www.cs.unc.edu/~weiwang/paper/DAMI04.pdf - Similar pages

#### [PDF] University of Texas at Arlington Dissertation Template

File Format: PDF/Adobe Acrobat - View as HTML

The algorithm defines the notion of anti-monotonicity in which ... The algorithm Episode

Discovery (ED) proposed in [1] is a data mining ...

itlab.uta.edu/ITLABWEB/ Students/sharma/theses/Bha05MS.pdf - Similar pages

#### [PDF] By Svetlozar Nestorov June 2000

File Format: PDF/Adobe Acrobat - View as HTML

data mining problems over relational data. The main features of query ocks are: ... that

obey this property are called monotone TUC. +. 98 or anti-monotone ...

www-db.stanford.edu/~evtimov/pubs/thesis.pdf - Similar pages

## [PDF] Towards Generic Pattern Mining

File Format: PDF/Adobe Acrobat - View as HTML

In this paper we propose the Data Mining Template Library, a collection ... patterns; in a

generic sense a pattern denotes links/relationships between sev- ...

www.cs.rpi.edu/~zaki/PS/ICFCA05 - Similar pages

#### <u>United States Patent Application: 0050027710</u>

... implementation based on the properties discussed above. ... the above-mentioned Apriori property on the ... are tightly coupled in the anti-monotonic relationships. ... appft1.uspto.gov/.../ - 71k - Supplemental Result - Cached - Similar pages

#### Extraction [CiteSeer; NEC Research Institute; Steve Lawrence, Kurt ...

analysis Local documents Information extraction Templates Display of ... data mining knowledge extraction information discovery data pattern ...

citeseer.ist.psu.edu/ InformationRetrieval/Extraction/hubs.html - 112k -Cached - Similar pages

# [PDF] FIRST WORKSHOP MULTI-RELATIONAL DATA MINING 2001 Editors: Arno ...

File Format: PDF/Adobe Acrobat - View as HTML

of Bayesian networks to relational domains. A PRM describes a template for a ... on patterns represented as Datalog queries is monotone with ... mrdm.dantec.nl/procmrdm.pdf - Similar pages

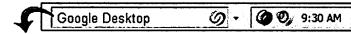
Try your search again on Google Book Search

Gooooooogle >

Result Page:

1 2 3 4 5 6 7 8 9 10

**Next** 



Free! Instantly find your email, files, media and web history. Download now.

anti monotonicity property discovery

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

Sign in



Web Images Groups News Froogle Local more » Advanced Search anti monotonicity property discovery relational Search

# Web Results 1 - 10 of about 479 for anti monotonicity property discovery relational data mining query tem

### [PDF] APA 5th Edition Template

File Format: PDF/Adobe Acrobat - View as HTML

Input: a mining guery with a preference, the preference is either monotonic or anti-

monotonic;. Output: the complete set of preference patterns;. Method: ...

www.cs.sfu.ca/~jpei/publications/prefmine-jdwm.pdf - Similar pages

# [PDF] Pushing Convertible Constraints in Frequent Itemset Mining

File Format: PDF/Adobe Acrobat - View as HTML

convertible anti-monotone wrt value descending order (see also Example 3). This property, provides great flexibility in data mining guery optimization. ...

www.cs.sfu.ca/~jpei/publications/conv jn.pdf - Similar pages

# [PDF] Simultaneous Optimization of Complex Mining Tasks with a ...

File Format: PDF/Adobe Acrobat - View as HTML

vision that a Knowledge Discovery and Data Mining System (KD-... issuing a sequence of

related complex mining queries, and 2) Multi-... www.cs.kent.edu/~jin/Papers/mgo.pdf - Similar pages

# [PDF] InfoMiner: Mining Surprising Periodic Patterns

File Format: PDF/Adobe Acrobat - View as HTML

Periodic pattern discovery is an important problem in mining time series data and has

wide ... (1) The surprise of a pattern occurrence is anti-monotonic ...

www.cs.unc.edu/~weiwang/paper/DAMI04.pdf - Similar pages

#### [PDF] University of Texas at Arlington Dissertation Template

File Format: PDF/Adobe Acrobat - View as HTML

The algorithm defines the notion of anti-monotonicity in which ... The algorithm Episode

Discovery (ED) proposed in [1] is a data mining ...

itlab.uta.edu/ITLABWEB/ Students/sharma/theses/Bha05MS.pdf - Similar pages

## [PDF] By Svetlozar Nestorov June 2000

File Format: PDF/Adobe Acrobat - View as HTML

data mining problems over relational data. The main features of query ocks are: ... that

obey this property are called monotone TUC. +. 98 or anti-monotone ...

www-db.stanford.edu/~evtimov/pubs/thesis.pdf - Similar pages

#### [PDF] Towards Generic Pattern Mining

File Format: PDF/Adobe Acrobat - View as HTML

In this paper we propose the Data Mining Template Library, a collection ... patterns; in a

generic sense a pattern denotes links/relationships between sev- ...

www.cs.rpi.edu/~zaki/PS/ICFCA05 - Similar pages

#### <u>United States Patent Application: 0050027710</u>

... implementation based on the properties discussed above. ... the above-mentioned Apriori property on the ... are tightly coupled in the anti-monotonic relationships. ... appft1.uspto.gov/.../ - 71k - Supplemental Result - Cached - Similar pages

#### Extraction [CiteSeer; NEC Research Institute; Steve Lawrence, Kurt ...

analysis Local documents Information extraction Templates Display of ... data mining knowledge extraction information discovery data pattern ...

citeseer.ist.psu.edu/ InformationRetrieval/Extraction/hubs.html - 112k -Cached - Similar pages

# [PDF] FIRST WORKSHOP MULTI-RELATIONAL DATA MINING 2001 Editors: Arno ...

File Format: PDF/Adobe Acrobat - View as HTML

of Bayesian networks to relational domains. A PRM describes a template for a ... on patterns represented as Datalog queries is monotone with ... mrdm.dantec.nl/procmrdm.pdf - Similar pages

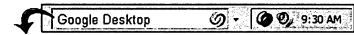
Try your search again on Google Book Search

Gooooooogle >

Result Page:

1 2 3 4 5 6 7 8 9 10

Next



Free! Instantly find your email, files, media and web history. Download now.

anti monotonicity property discovery

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

Sign in



 Web
 Images
 Groups
 News
 Froogle
 Local
 more »

 anti monotonicity mining attribute relational dal
 Search
 Advanced Search Preferences

# Web Results 11 - 20 of about 10,400 for anti monotonicity mining attribute relational data set. (0.29 second

# [PDF] Mining Multi-Dimensional Constrained Gradients in Data Cubes

File Format: PDF/Adobe Acrobat - View as HTML

on dimensional attributes confines the set of probe cells ... Anti-monotonicity is very

useful for pruning. It states that if a cell ...

www-sal.cs.uiuc.edu/~hanj/pdf/vldb01.pdf - Similar pages

# [PDF] Data Mining: Concepts and Techniques

File Format: PDF/Adobe Acrobat - View as HTML

frequent set computation. January 17, 2001. Data Mining: Concepts and Techniques. 71.

Characterization of. Anti-Monotonicity Constraints ...

www.ir.iit.edu/~dagr/DataMiningCourse/ Spring2001/BookNotes/6asso.pdf - Similar pages

# [PDF] Inductive Databases: Towards a New Generation of Databases For ...

File Format: PDF/Adobe Acrobat

That is, the generic data mining task is to find the set Th ... rithms make use of the

monotonicity or anti-monotonicity. of such constraints, as in [5] and ...

dx.doi.org/10.1109/DEXA.2005.116 - Similar pages

# ACM: Digital Library: Communications of the ACM

In the first group, identification of the **anti-monotonicity** property that all subsets of a ... The input **data set** to the **data mining** process is not static, ... dx.doi.org/10.1145/545151.545176 - <u>Similar pages</u>

## [PDF] Mining Association Rules With Negation In Inductive Databases 1 ...

File Format: PDF/Adobe Acrobat - View as HTML

candidate generation relies on the anti-monotonicity of the frequency constraint ... to mine

all null rules in the original dataset; we just treat the new ...

www.cs.indiana.edu/~gefletch/papers/negation.pdf - Similar pages

# <sub>[РФF]</sub> <u>Visualization to assist users in association rule **mining** tasks</u>

File Format: PDF/Adobe Acrobat - View as HTML

swer, due to the anti-monotonicity property of the support. ... Set-oriented mining for

association rules. in relational databases. In Proc. Int'l. ...

www.sbbd-sbes2005.ufu.br/arquivos/11866.pdf - Similar pages

# [PPT] CS490D: Introduction to Data Mining Chris Clifton

File Format: Microsoft Powerpoint 97 - View as HTML

Key point: find a function which is weaker but ensures certain anti-monotonicity ... Data

cube: A relational aggregation operator generalizing group-by, ...

www.cs.purdue.edu/homes/clifton/cs490d/DW.ppt - Similar pages

### [PDF] Chapter 2 CONSTRAINT-BASED DATA MINING

File Format: PDF/Adobe Acrobat - View as HTML

itemset mining (FIM) where the data is a set of transactions, the patterns are ... the

specialization relation has an anti-monotonicity property. ...

liris.cnrs.fr/~jboulica/constraints\_handbook.pdf - Similar pages

## [PS] Incremental execution of the MINE RULE operator Marco Botta and ...

File Format: Adobe PostScript - View as HTML

be two mining queries, extracting from the same source data rule sets R ... exploit the

principles of succinctness and **anti-monotonicity** of the constraints ... www.di.unito.it/~meo/Pubblicazioni/RT66-2002.ps - <u>Similar pages</u>

[PPT] Tutorial Inductive Databases

File Format: Microsoft Powerpoint 97 - View as HTML

This seminar focuses on database aspects of data mining ... Anti-monotonicity of q wrt .

q is anti-monotone wrt if and only if ...

www.informatik.uni-freiburg.de/ ~ml/teaching/ss04/dm/dmdbseminar.ppt - Similar pages

d Gooooooooogle ▶

Result Page: **Previous** 1 2 3 4 5 6 7 8 9 1011 **Next** 

anti monotonicity mining attribute rel

Search within results | Language Tools | Search Tips

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

Sign in



web Images Groups News Froogle Local more >>

anti monotonicity mining attribute relational dat

Search

Advanced Search

Preferences

Web Results 21 - 30 of about 10,400 for anti-monotonicity mining attribute relational data set. (0.38 second

### IPDFI TSET: AN ALGORITHM FOR MINING FREQUENT TEMPORAL PATTERNS

File Format: PDF/Adobe Acrobat - View as HTML mining is to determine in the dataset all the frequent sequences whose support ... This property, also called anti-monotonicity property, ... www.lsi.us.es/~aguilar/ecml2004/FP7.PDF - Similar pages

# [PS] Database Mining Integration Seminar Report Submitted in partial ...

File Format: Adobe PostScript - View as HTML

... the data mining utilities must assume a relational backend. ... However due to antimonotonicity we test the sets for satisfaction of the constraints ... www.prism.gatech.edu/~gtg463i/publications/seminar.ps - <u>Similar pages</u>

# [PDF] Mining Open Source Software(OSS) Data using Association Rules Network

File Format: PDF/Adobe Acrobat - View as HTML

tween **relational data attributes** in general and OSS **data** in particular. ... tage of the **anti-monotonicity** property exhibited by the support level: If ... www.it.usyd.edu.au/~chawla/ publications/chawlapakdd03.pdf - <u>Similar pages</u>

# [PDF] Monotone Constraints in Frequent Tree Mining

File Format: PDF/Adobe Acrobat - View as HTML

iments on synthetic **data sets**, to evaluate the mono-. tone tree **mining** algorithm. ... fined over a tree t has the **anti-monotonicity** property ... www.cs.uu.nl/people/jknijf/benelearn2005.pdf - <u>Similar pages</u>

#### [PDF] Constraint Based Mining of First Order Sequences in SeqLog

File Format: PDF/Adobe Acrobat - View as HTML

from these atoms because of the **anti-monotonicity**—a contradiction. ... frequent query discovery in multi-**relational data mining** and inductive logic pro- ... www.informatik.uni-freiburg.de/ ~ml/papers/conf-kdid-LeeR02.pdf - <u>Similar pages</u> [ More results from www.informatik.uni-freiburg.de ]

# [PDF] On Monotone Data Mining Languages

File Format: PDF/Adobe Acrobat - View as HTML

propose a logic, called **Data Mining** Logic (DML), that extends **relational** tuple. calculus with variables ranging over **attributes** and over **sets** of **attributes** ... www.adrem.ua.ac.be/bibrem/pubs/CALDERSDBPL01.pdf - Similar pages

# [PDF] Mining alpha/beta concepts as relevant bi-sets from transactional data

File Format: PDF/Adobe Acrobat - View as HTML

this paper, we consider bi-set mining from transactional data. More precisely, we ... is anti-monotonic wrt our specialization relation on bi-sets. ... liris.cnrs.fr/~jboulica/kdid04.pdf - Similar pages

[ More results from liris.cnrs.fr ]

# [PDF] eurandom.pdf — April 29, 2003 — 1 • Patterns Relational Data ...

File Format: PDF/Adobe Acrobat - View as HTML

the goal of **relational data mining**. Generalise **Data Mining** to Multiple Tables ... on a **set** of **anti-**HIV active compounds, substructures ...

www.eurandom.tue.nl/workshops/ 2003/datamining/siebes.pdf - Similar pages

# [PDF] Mining Quantitative Frequent Itemsets Using Adaptive Density-based ...

File Format: PDF/Adobe Acrobat

levelwise algorithm with the (anti-)monotonicity property. of the dense clusters. ...

attribute space. The similar data sets were used to evalu- ...

doi.ieeecomputersociety.org/10.1109/ICDM.2005.100 - Similar pages

### грет No Slide Title

File Format: Microsoft Powerpoint 97 - View as HTML

Data Mining: Concepts and Techniques. 70. Property of Constraints: Anti-Monotone. Anti-

monotonicity: If a set S violates the constraint, any superset of S ...

www.cs.sfu.ca/~han/bk/6asso.ppt - Similar pages

# ◆ Gooooooooogle ▶

Result Page: **Previous** 1 2 3 4 5 6 7 8 9 101112

anti monotonicity mining attribute rel Search



Search within results | Language Tools | Search Tips

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

<u>Sign in</u>



Web Images Groups News Froogle Local Advanced Search mining template mining attributes relatinal data Search

Web Results 1 - 10 of about 14 for mining template mining attributes relatinal data set . (0.46 seconds)

Did you mean: mining template mining attributes *relational* data set

ppp Discovering database summaries through refinements of fuzzy ... File Format: PDF/Adobe Acrobat - View as HTML one of the major parts of Data. mining techniques, which provide the user with com- ... hypothesis template specifies which attributes of the Data ... 140.98.193.112/iel2/945/7004/00283034. pdf?tp=&isnumber=7004&arnumber=283034 -Supplemental Result - Similar pages

[PDF] Discovering database summaries through refinements of fuzzy ... File Format: PDF/Adobe Acrobat one of the major parts of data. mining techniques, which provide the user with com-... hypothesis template specifies which attributes of the data ... ieeexplore.ieee.org/iel2/ 945/7004/00283034.pdf?arnumber=283034 - Similar pages

# Tier Questions Page

... Mine won't work ... 286475 XSD, validating Attributes values and elements. 286476 How to export excel ... 286599 Making a template for a reply rule in MS Outlook 2003. ... www.experts-exchange.com/T 5728.html - 58k - Supplemental Result -Cached - Similar pages

System and method for the creation and automatic deployment of ... MSTR035— Relatinal OLAP Interface, Programmer's Reference and SDK Guide, ... Patterns and Trends by Applying OLAP and Data Mining Technology on Web Logs", ... www.freepatentsonline.com/6885734.html - 212k - Cached - Similar pages

Extending the database relational model to capture more meaning Carlos Ordonez, Horizontal aggregations for building tabular data sets, Proceedings of the 9th ACM SIGMOD workshop on Research issues in data mining and ... portal.acm.org/citation.cfm?id=320107.320109 - Similar pages

#### XML Quotes in 2001

I see RDF and tend to think of data mining and discovery. ... When we have the notion of defaulted or fixed attributes, then an XML document cannot truly ... www.ibiblio.org/xml/quotes2001.html - 190k - Cached - Similar pages

#### Category Browsing Results

Data MINING and Uncertain Reasoning: An Integrated Approach ... Each CD-ROM provides powerful DESIGN templates discussed In the Books In A ready-TO-Use ... technical-books.com/cf/browse.cfm?Main=COMPUTERS%2C%20PC&Sub1=DATABASES - 179k - Supplemental Result - Cached - Similar pages

ftp.linux.cz/pub/bibliography/Database/time.bib.gz?N=A

File Format: Unrecognized - View as HTML Supplemental Result - Similar pages

#### OJB

One is a data mining toolset, and the other is a massive National Science ... In the fielddescriptor of this attribute set the attribute locking="true". ... db.apache.org/ojb.1\_0\_2/site.html - 1130k - Cached - Similar pages

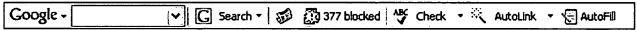
PDFJ Table of contents

File Format: PDF/Adobe Acrobat - View as HTML and about 20 or 25 universities and research groups like mine. ... Java attribute for it. In the field-descriptor of this attribute set the attribute ... db.apache.org/ojb.1\_0\_2/site.pdf - Similar pages

Did you mean to search for: mining template mining attributes relational data set

Google ▶ Result Page:

Free! Get the Google Toolbar. Download Now - About Toolbar



mining template mining attributes rel

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

Sign in



 Web
 Images
 Groups
 News
 Froogle
 Local
 more »

 pruning procedure pattern template
 Search
 Advanced Search Preferences

Web

Results 1 - 10 of about 132,000 for pruning procedure pattern template. (0.39 seconds)

Did you mean: planting procedure pattern template

Sponsored Links

**Computer Vision** 

The standard belief propagation algorithm is augmented by a **pruning procedure** and a novel technique, inspired by the 20 Questions (divide-and-conquer) ...

www.ski.org/Rehab/ComputerVision/James/ - 15k - Cached - Similar pages

Gardening Software
Gardening software.
Free to use. Get info instantly.
www.puredirectory.com

[Paper] Cluttered Scene Understanding and Augmentation

If a cluster contains more fea- tures than the **template**, the cluster needs to be ... To recognize free-form object, the **pruning procedure** as well as ... www.actapress.com/PDFViewer.aspx?paperId=18830 - Similar pages

Citations: Interpretation of artificial neural networks: Mapping ...

Threshold **pruning** with threshold = 0:01 is applied, followed by the rule and ... **Template**-Based **Procedures** for Neural Network Interpretation - Alexander, ... citeseer.ist.psu.edu/context/75128/0 - 37k - <u>Cached</u> - <u>Similar pages</u>

[PS] Error-Driven Pruning of Treebank Grammars for Base Noun Phrase ...

File Format: Adobe PostScript - View as HTML

Our general **pruning procedure** is shown in Figure 3. ... The treebank approach and the R&M approach without lexical **templates** are shown to perform comparably ... www.cs.cornell.edu/home/cardie/papers/acl98.ps - <u>Similar pages</u>

[PS] Combining Error-Driven Pruning and Classification for Partial ...

File Format: Adobe PostScript - View as HTML

The **patterns**, or grammars, that drive each stage ofprocessing ... 3.1 Error-Driven **Pruning** Empire's grammar **pruning procedure** is shown in Fig-ure 3. ... www.cs.cornell.edu/home/cardie/papers/icml-99.ps - Similar pages

# [PDF] THE GRAPHICAL MODELS TOOLKIT: AN OPEN SOURCE SOFTWARE SYSTEM FOR ...

File Format: PDF/Adobe Acrobat - View as HTML

Although the statistical approach to **pattern** classification has been ... bined with GMTK's beam-**pruning procedure** for further memory. savings. ... ssli.ee.washington.edu/people/ bilmes/mypapers/gmtk\_icassp02.pdf - Similar pages

[PDF] Analysis of a complexity based pruning scheme for classification trees

File Format: PDF/Adobe Acrobat - View as HTML

A comprehensive treatment of probabilistic **pattern** recognition ... The CART **pruning procedure** selects a subtree of the initial tree that minimizes a ... www.stat.unc.edu/faculty/nobel/links/prune.pdf - Similar pages

[PDF] Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA

File Format: PDF/Adobe Acrobat - View as HTML

deformable **template** approach to match handwritten digits. and 3D objects. ... The end result of our fast **pruning procedure** is a set of ...

www.cs.sfu.ca/~mori/research/papers/mori\_cvpr03.pdf - Similar pages

#### [PDF] Shape-Based Recognition of Wiry Objects

File Format: PDF/Adobe Acrobat - View as HTML

while template-based techniques are effective for some objects, we feel it is worth ...

Specifically, our pruning procedure described in Section III-D ...

www.cs.cmu.edu/~owenc/carmichael\_hebert\_pami.pdf - Similar pages

#### [PDF] Matching and Retrieving Sequential Patterns Under Regression

File Format: PDF/Adobe Acrobat - View as HTML

The **procedure** above combines of the SLR-**pruning** and ... How about the power of SLR+FFT-**pruning?** • Does the linear **pattern** retrieval make sense in the real ...

www.cedar.buffalo.edu/~govind/regression.pdf - Similar pages

Did you mean to search for: *planting* procedure pattern template

Try your search again on Google Book Search

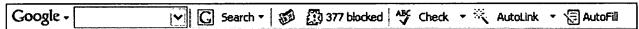
# Goooooooogle >

Result Page:

1 2 3 4 5 6 7 8 9 10

Next

Free! Get the Google Toolbar. Download Now - About Toolbar



pruning procedure pattern template

Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

Sign in



 Web
 Images
 Groups
 News
 Froogle
 Local
 more »

 query template pruning procedure pattern temi
 Search
 Advanced Search

 Preferences

Web Results 11 - 20 of about 55,000 for query template pruning procedure pattern template. (0.10 seconds

#### IPDFI AAAI Proceedings Template

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> employs a weight sampling **procedure** to finally select particularly ... Not **pruning** and using the Laplace correction had been shown to improve the ... pages.stern.nyu.edu/~fprovost/ Papers/maytal-is-01-03.pdf - <u>Similar pages</u>

#### [Paper] Architectural Pattern Detection - A Case Study

The **pattern templates** are stored in a **template** library and the stored ... **Pruning** the branches created during search using Prolog's "cut" or "if then else" ... www.actapress.com/PDFViewer.aspx?paperId=22414 - <u>Similar pages</u>

## [DOC] Proceedings Template - WORD

File Format: Microsoft Word - <u>View as HTML</u>
We detail the **procedures** used to create the Omega base ontology and to ... Further investigation of shishkebab **patterns** within the growing ontology are part ... www.isi.edu/~philpot/papers/ilc2003/ilc2003final.doc - <u>Similar pages</u>

## [PDF] Automatic Ordering of Subgoals - a Machine Learning Approach

File Format: PDF/Adobe Acrobat - <u>View as HTML</u>
LASSY currently uses a weighted set of **query templates** as its task model. ... For each **query**, the **procedure** checks to see whether its calling **pattern** ... www.cs.technion.ac.il/~shaulm/ papers/pdf/Markovitch-Scott-naclp1989.pdf - <u>Similar pages</u>

#### [PDF] Proceedings Template - WORD

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> **template** with data from a database. HTML forms modify the data ... **pruning**, but is limited to binary trees of depth at most five. The ...
www.cs.cmu.edu/~einat/iiweb-2004.pdf - <u>Similar pages</u>

#### [DOC] Proceedings Template - WORD

File Format: Microsoft Word - View as HTML

The state-**pruning** techniques described by Deshpanse and Karypis [7] offer an approach ... **Patterns** of search: Analyzing and modeling web **query** refinement. ... sifaka.cs.uiuc.edu/xshen/research\_files/ MSRTechReport\_topicdynamics.doc - Similar pages

## Papakonstantinou, Y.; Gupta, A.; Haas, L.: Capabilities-Based ...

2 8 3.1 **Pruning** \Represented" CSQs A large number of unneeded CSQs are generated by **templates** ... Answering **queries** using **templates** with binding **patterns**. ... dbpubs.stanford.edu/pub/1995-2 - 82k - <u>Cached</u> - <u>Similar pages</u>

#### [PDF] Proceedings Template - WORD

File Format: PDF/Adobe Acrobat - <u>View as HTML</u>
Traditional **query** optimizers rely on the accuracy of estimated ... plans that remain after **pruning** are the three plans corresponding. to the minimum C ... www.cs.wisc.edu/~pedro/papers/sigmod2005\_rio.pdf - <u>Similar pages</u>

## [PDF] Recognizing Cars Abstract 1 Introduction

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> isting LPR systems is correlation-based **template** match- ... lowing keypoint **pruning** 

procedures:. • Limit horizontal distance between matching keypoints. ... www-cse.ucsd.edu/classes/ fa05/cse252c/dlagnekov\_belongie.pdf - Similar pages

[PDF] Finding and Clustering Database User Sessions

File Format: PDF/Adobe Acrobat - View as HTML

The discovered patterns can be used to predict incoming user queries based ... template.

The query template represents a set of queries that have the ...

www.yorku.ca/jhuang/paper/dasfaa05.pdf - Similar pages



Result Page: **Previous** 1 2 3 4 5 6 7 8 9 1011

query template pruning procedure p

Search within results | Language Tools | Search Tips

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

Sign in



Web Images Groups News Froogle Local more »

discovery relational data mining query templaty

Search

Preferences

### Web Results 1 - 10 of about 19,400 for discovery relational data mining guery template pruning procedure

### [PDF] Contents

File Format: PDF/Adobe Acrobat - <u>View as HTML</u>
We now focus on specific **data mining** and **pattern discovery** tasks that are often ... straints on **pattern templates**, such as wildcards with user-specified ... maya.cs.depaul.edu/~mobasher/papers/IC-Handbook-04.pdf - <u>Similar pages</u>

### грог Untangling Text Data Mining

File Format: PDF/Adobe Acrobat - View as HTML of discovery of trends and patterns within tex-. tual data for more general purpose usage. ... port for issuing sequences of queries and related ... www.sims.berkeley.edu/~hearst/papers/acl99.pdf - Similar pages

## **Untangling Text Data Mining**

Instead, data mining applications tend to be (semi)automated discovery of trends ... sequences of queries and related operations across text collections, ... www.sims.berkeley.edu/ ~hearst/papers/acl99/acl99-tdm.html - 40k - Cached - Similar pages

### [PDF] Querying Multiple Sets of Discovered Rules

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> **query** languages. One of the earliest **data mining query** language. is the one based on **templates** [15]. In this technique, the user uses. a **template** to specify ... www.cs.uic.edu/~liub/publications/kdd2002-RuleQL.pdf - <u>Similar pages</u>

# [РРБ] An Indexing Scheme for Fast Similarity Search in Large Time Series ... File Format: PDF/Adobe Acrobat - View as HTML

We begin by placing a **template** grid over the **query** Q. This **template** is exactly the same as the one ... Knowledge **Discovery** and **Data Mining**. pp 16-22, AAAI ... www.ics.uci.edu/~pazzani/Publications/ssdb99.pdf - Similar pages

#### [PDF] APA 5th Edition Template

File Format: PDF/Adobe Acrobat - View as HTML

The research on preference-based frequent **pattern mining** is **related** to two categories of. previous studies: preference **queries** and frequent **pattern mining** ... www.cs.sfu.ca/~jpei/publications/prefmine-jdwm.pdf - <u>Similar pages</u>

#### [PDF] Final Report

File Format: PDF/Adobe Acrobat - View as HTML

David J. Hand: **Pattern** detection and **discovery**. **Data mining** comprises ... background knowledge and us a **template**-based **mining** language to exploit it during ... www.esf.org/generic/1362/0110PReport.pdf - <u>Similar pages</u>

# Mining Software Engineering Data: A Survey - Tool Descriptions

AIRA includes different **templates** that can be used, one based on frames, ... Provides **Discovery** of **Data Patterns** - Yes **Data Mining** in CrossTables Analysis ... www.dacs.dtic.mil/techs/datamining/descriptions.shtml - 184k - <u>Cached</u> - <u>Similar pages</u>

## Citations: Finding Interesting Rules from Large Sets of Discovered ...

... of data mining to discover interesting relationships among attributes in ... Templates [10], pattern expressions that describe the form of rules that ...

discovery relational data mining query template pruning procedure pattern template - Goo... Page 2 of 2

citeseer.ist.psu.edu/context/35674/0 - 35k - Cached - Similar pages

[PDF] Inspec Alerts

File Format: PDF/Adobe Acrobat - <u>View as HTML</u> sizeable set of **pattern templates** to be generated for **mining** or. **query**. ... Integrating **query** processing and **data mining** in. **relational** DBMSs ... inspec.iee.org/inspec/currentawareness/ pdfs/products/samples/IEE\_data\_mining.pdf - Similar pages

Try your search again on Google Book Search

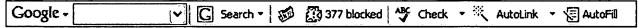
Gooooooogle >

Result Page: 1

1 2 3 4 5 6 7 8 9 10 No

<u>Next</u>

Free! Get the Google Toolbar. <u>Download Now</u> - <u>About Toolbar</u>



discovery relational data mining que

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: 

The ACM Digital Library C The Guide

mining templates

SEARCH



Feedback Report a problem Satisfaction survey

Terms used mining templates

Found 10,524 of 171,143

Sort results by

 $\Diamond$ relevance

Save results to a Binder ? Search Tips

Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Display results

expanded form  $\nabla$ 

window

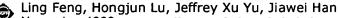
Result page: 1 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Results 1 - 20 of 200

Relevance scale

Mining inter-transaction associations with templates



November 1999 Proceedings of the eighth international conference on Information and knowledge management

**Publisher: ACM Press** 

Full text available: ndf(1.13 MB)

Additional Information: full citation, abstract, references, citings, index terms

Multi-dimensional, inter-transaction association rules extend the traditional association rules to describe more general associations among items with multiple properties cross transactions. "After McDonald and Burger King open branches, KFC will open a branch two months later and one mile away" is an example of such rules. Since the number of potential inter-transaction association rules tends to be extremely large, mining intertransaction associations poses more chall ...

<sup>2</sup> A template model for multidimensional inter-transactional association rules Ling Feng, Jeffrey Xu Yu, Hongjun Lu, Jiawei Han

October 2002 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 11 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(313.62 KB) Additional Information: full citation, abstract, index terms

Multidimensional inter-transactional association rules extend the traditional association rules to describe more general associations among items with multiple properties across transactions. "After McDonald and Burger King open branches, KFC will open a branch two months later and one mile away" is an example of such rules. Since the number of potential inter-transactional association rules tends to be extremely large, mining intertransactional associations poses more challe ...

Keywords: Intra-transactional/inter-transactional association rules, Multidimensional context, Template model

Research track: Eliminating noisy information in Web pages for data mining

Lan Yi, Bing Liu, Xiaoli Li

August 2003 Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

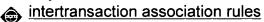


Full text available: pdf(468.87 KB) Additional Information: full citation, abstract, references, citings, index

A commercial Web page typically contains many information blocks. Apart from the main content blocks, it usually has such blocks as navigation panels, copyright and privacy notices, and advertisements (for business purposes and for easy user access). We call these blocks that are not the main content blocks of the page the noisy blocks. We show that the information contained in these noisy blocks can seriously harm Web data mining. Eliminating these noises is thus of great importance. In this pa ...

**Keywords**: Web mining, noise detection, noise elimination

4 Beyond intratransaction association analysis: mining multidimensional



Hongjun Lu, Ling Feng, Jiawei Han

October 2000 ACM Transactions on Information Systems (TOIS), Volume 18 Issue 4

Publisher: ACM Press

Full text available: pdf(1.31 MB)

Additional Information: full citation, abstract, references, citings, index terms

In this paper, we extend the scope of mining association rules from traditional singledimensional intratransaction associations, to multidimensional intertransaction associations. Intratransaction associations are the associations among items with the same transaction, where the notion of the transaction could be the items bought by the same customer, the events happened on the same day, and so on. However, an intertransaction association ...

**Keywords:** association rules, data mining, intra/intertransaction, multidimensional context

5 Industrial and practical experience track paper session 1: The volume and evolution



of web page templates

David Gibson, Kunal Punera, Andrew Tomkins

May 2005 Special interest tracks and posters of the 14th international conference on World Wide Web

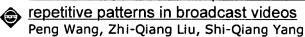
Publisher: ACM Press

Full text available: 🔂 pdf(249.32 KB) Additional Information: full citation, abstract, references, index terms

Web pages contain a combination of unique content and template material, which is present across multiple pages and used primarily for formatting, navigation, and branding. We study the nature, evolution, and prevalence of these templates on the web. As part of this work, we develop new randomized algorithms for template extraction that perform approximately twenty times faster than existing approaches with similar quality. Our results show that 40--50% of the content on the web is templa ...

Keywords: algorithms, boilerplate, data cleaning, data mining, templates, web mining

6 Poster 3: content track: A probabilistic template-based approach to discovering



November 2005 Proceedings of the 13th annual ACM international conference on **Multimedia MULTIMEDIA '05** 

Publisher: ACM Press

Full text available: pdf(163.58 KB) Additional Information: full citation, abstract, references, index terms

There are usually repetitive sub-segments in broadcast videos, which may be associated with high-level concepts or events, e.g., news footage, repeated scores in basketball. Unsupervised mining techniques provide generic solutions to discovering such temporal patterns in various video genres, which are currently the subject of great interests to researchers working on multimedia content analysis. In this paper, we propose a novel approach to automatically detecting repetitive patterns in a video ...

Keywords: probabilistic template, repetitive pattern discovery, video mining

7 Description and Analysis: Template detection via data mining and its applications Ziv Bar-Yossef, Sridhar Rajagopalan



May 2002 Proceedings of the 11th international conference on World Wide Web

Publisher: ACM Press Additional Information: full citation, abstract, references, citings, index Full text available: pdf(404.46 KB) terms

We formulate and propose the template detection problem, and suggest a practical solution for it based on counting frequent item sets. We show that the use of templates is pervasive on the web. We describe three principles, which characterize the assumptions made by hypertext information retrieval (IR) and data mining (DM) systems, and show that templates are a major source of violation of these principles. As a consequence, basic "pure" implementations of simple search algorithms coupled with t ...

**Keywords**: data mining, hypertext, information retrieval, web searching

Analysis of navigation behaviour in web sites integrating multiple information systems Bettina Berendt, Myra Spiliopoulou



March 2000 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 9 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(281.14 KB) Additional Information: full citation, abstract, citings, index terms

The analysis of web usage has mostly focused on sites composed of conventional static pages. However, huge amounts of information available in the web come from databases or other data collections and are presented to the users in the form of dynamically generated pages. The query interfaces of such sites allow the specification of many search criteria. Their generated results support navigation to pages of results combining crosslinked data from many sources. For the analysis of visitor naviga ...

Keywords: Conceptual hierarchies, Data mining, Query capabilities, Web databases, Web query interfaces, Web usage mining

9 Extended abstract: Database intrusion detection based on user query frequent



itemsets mining with item constraints

Yong Zhong, Xiao-lin Qin

November 2004 Proceedings of the 3rd international conference on Information security InfoSecu '04

Publisher: ACM Press

Full text available: pdf(176.00 KB) Additional Information: full citation, abstract, references, index terms

The paper presents a database intrusion detection algorithm based on user query frequent itemsets with item constraints. Firstly, the paper discusses the method to mine database user query frequent itemsets by using query template. Secondly, the paper discusses the constrained query templates that are used to reduce the number of frequent itemsets and improve system performance. At last, the paper discusses the algorithm design and its application example.

Keywords: data mining, database security, frequent itemsets, intrusion detection

10 Background for association rules and cost estimate of selected mining algorithms Jia Liang Han, Ashley W. Plank



November 1996 Proceedings of the fifth international conference on Information and knowledge management

Publisher: ACM Press

Full text available: 📆 pdf(928.44 KB) Additional Information: full citation, references, citings, index terms

Keywords: algorithm analysis, association rules, background, data mining, memory management, optimization, scalability

11 An object-oriented approach to multi-level association rule mining





Scott Fortin, Ling Liu

November 1996 Proceedings of the fifth international conference on Information and knowledge management

Publisher: ACM Press

Full text available: 📆 pdf(996.66 KB) Additional Information: full citation, references, citings, index terms

12 <u>Bug localization: PR-Miner: automatically extracting implicit programming rules and</u>





detecting violations in large software code Zhenmin Li, Yuanyuan Zhou

September 2005 Proceedings of the 10th European software engineering conference held jointly with 13th ACM SIGSOFT international symposium on Foundations of software engineering ESEC/FSE-13

Publisher: ACM Press

Full text available: pdf(228.58 KB) Additional Information: full citation, abstract, references, index terms

Programs usually follow many implicit programming rules, most of which are too tedious to be documented by programmers. When these rules are violated by programmers who are unaware of or forget about them, defects can be easily introduced. Therefore, it is highly desirable to have tools to automatically extract such rules and also to automatically detect violations. Previous work in this direction focuses on simple function-pair based programming rules and additionally requires programmer ...

**Keywords:** automated specification generation, automated violation detection, data mining for software engineering, pattern recognition, programming rules, static analysis

13 Editorial: special issue on web content mining



Bing Liu, Kevin Chen-Chuan-Chang

December 2004 ACM SIGKDD Explorations Newsletter, Volume 6 Issue 2

Publisher: ACM Press

Full text available: The pdf(178.32 KB) Additional Information: full citation, abstract, references

With the phenomenal growth of the Web, there is an everincreasing volume of data and information published in numerous Web pages. The research in Web mining aims to develop new techniques to effectively extract and mine useful knowledge or information from these Web pages [8]. Due to the heterogeneity and lack of structure of Web data, automated discovery of targeted or unexpected knowledge/information is a challenging task. It calls for novel methods that draw from a wide range of fields spanni ...

14 Intelligent jurisprudence research: a new concept



Rosina Weber

June 1999 Proceedings of the 7th international conference on Artificial intelligence and law

**Publisher: ACM Press** 

Full text available: pdf(1.20 MB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

Intelligent Jurisprudence Research (IJR) is a concept that consists in performing jurisprudence research with a computational tool that employs Artificial Intelligence (AI) techniques. Jurisprudence research is the search employed by judicial professionals when seeking for past legal situations that may be useful to a legal activity. When humans perform jurisprudence research, they employ analogical reasoning in comparing a given actual situation with past decisions, noting the affinities b ...

**Keywords**: artificial intelligence and law, case-based reasoning, case-based retrieval, jurisprudence

15 Exploratory mining and pruning optimizations of constrained associations rules



Raymond T. Ng, Laks V. S. Lakshmanan, Jiawei Han, Alex Pang

June 1998 ACM SIGMOD Record, Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98, Volume 27 Issue 2

**Publisher: ACM Press** 

Full text available: pdf(1.65 MB)

Additional Information: full citation, abstract, references, citings, index terms

From the standpoint of supporting human-centered discovery of knowledge, the presentday model of mining association rules suffers from the following serious shortcomings: (i) lack of user exploration and control, (ii) lack of focus, and (iii) rigid notion of relationships. In effect, this model functions as a black-box, admitting little user interaction in between. We propose, in this paper, an architecture that opens up the black-box, and supports constraint-based, human-centered explorat ...

16 Towards on-line analytical mining in large databases



Jiawei Han

March 1998 ACM SIGMOD Record, Volume 27 Issue 1

**Publisher: ACM Press** 

Full text available: Topdf(387.04 KB) Additional Information: full citation, abstract, citings, index terms

Great efforts have been paid in the Intelligent Database Systems Research Lab for the research and development of efficient data mining methods and construction of on-line analytical data mining systems. Our work has been focused on the integration of data mining and OLAP technologies and the development of scalable, integrated, and multiple data mining functions. A data mining system, DBMiner, has been developed for interactive mining of multiple-level knowledge in large relational databases and ...

17 Industrial/government track: Capturing best practice for microarray gene expression



data analysis

Gregory Piatetsky-Shapiro, Tom Khabaza, Sridhar Ramaswamy

August 2003 Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: Topdf(874.16 KB) Additional Information: full citation, abstract, references, index terms

Analyzing gene expression data from microarray devices has many important application in medicine and biology, but presents significant challenges to data mining. Microarray data typically has many attributes (genes) and few examples (samples), making the process of correctly analyzing such data difficult to formulate and prone to common mistakes. For this reason it is unusually important to capture and record good practices for this form of data mining. This paper presents a process for analyzi ...

Keywords: Clementine, application template, data mining process, gene expression, microarrays

18 Poster papers - short papers: Extracting unstructured data from template generated





web documents

Ling Ma, Nazli Goharian, Abdur Chowdhury, Misun Chung

November 2003 Proceedings of the twelfth international conference on Information and knowledge management

**Publisher: ACM Press** 

Full text available: pdf(210.48 KB)

Additional Information: full citation, abstract, references, citings, index terms

We propose a novel approach that identifies web page templates and extracts the unstructured data. Extracting only the body of the page and eliminating the template increases the retrieval precision for the queries that generate irrelevant results. We believe that by reducing the number of irrelevant results; the users are encouraged to go back to a given site to search. Our experimental results on several different web sites and on the whole cnnfn collection demonstrate the feasibility of our a ...

**Keywords**: automatic template removal, information retrieval, retrieval accuracy, text extraction

19 Research track poster: Simultaneous optimization of complex mining tasks with a





knowledgeable cache

Ruoming Jin, Kaushik Sinha, Gagan Agrawal

August 2005 Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05

Publisher: ACM Press

Full text available: 🔂 pdf(750.91 KB) Additional Information: full citation, abstract, references, index terms

With an increasing use of data mining tools and techniques, we envision that a Knowledge Discovery and Data Mining System (KDDMS) will have to support and optimize for the following scenarios: 1) Sequence of Queries: A user may analyze one or more datasets by issuing a sequence of related complex mining queries, and 2) Multiple Simultaneous Queries: Several users may be analyzing a set of datasets concurrently, and may issue related complex queries. This paper presents a systematic ...

Keywords: frequent pattern mining, knowledgeable cache, multiple query optimization

20 Bug localization: DynaMine: finding common error patterns by mining software



revision histories

Benjamin Livshits, Thomas Zimmermann

September 2005 Proceedings of the 10th European software engineering conference held jointly with 13th ACM SIGSOFT international symposium on Foundations of software engineering ESEC/FSE-13

**Publisher: ACM Press** 

Full text available: pdf(182.92 KB) Additional Information: full citation, abstract, references, index terms

A great deal of attention has lately been given to addressing software bugs such as errors in operating system drivers or security bugs. However, there are many other lesser known errors specific to individual applications or APIs and these violations of applicationspecific coding rules are responsible for a multitude of errors. In this paper we propose DynaMine, a tool that analyzes source code check-ins to find highly correlated method calls as well as common bug fixes in order to automatical ...

Keywords: coding patterns, data mining, dynamic analysis, error patterns, one-line check-ins, revision histories, software bugs

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library O The Guide

merge join sort pattern templates

SEARCH



Feedback Report a problem Satisfaction survey

Terms used merge join sort pattern templates

Found 33,487 of 171,143

Sort results

relevance  $\Box$ by

Save results to a Binder Search Tips Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Display results

expanded form  $\nabla$ 

window

Result page: 1 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Results 1 - 20 of 200

Relevance scale

1 Design patterns for sorting



Dung "Zung" Nguyen, Stephen B. Wong

February 2001 ACM SIGCSE Bulletin, Proceedings of the thirty-second SIGCSE technical symposium on Computer Science Education SIGCSE '01,

Volume 33 Issue 1

**Publisher: ACM Press** 

Full text available: pdf(78.20 KB)

Additional Information: full citation, abstract, references, citings, index terms

Drawing on Merritt's divide-and-conquer sorting taxonomy [1], we model comparisonbased sorting as an abstract class with a template method to perform the sort by relegating the splitting and joining of arrays to its concrete subclasses. Comparison on objects is carried out via an abstract ordering strategy. This reduces code complexity and simplifies the analyses of the various concrete sorting algorithms. Performance measurements and visualizations can be added without modifying any code by ut ...

<u>Fast detection of communication patterns in distributed executions</u>

Thomas Kunz, Michiel F. H. Seuren

November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research

Publisher: IBM Press

Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

3 Query evaluation techniques for large databases



Goetz Graefe

June 1993 ACM Computing Surveys (CSUR), Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(9.37 MB)

Additional Information: full citation, abstract, references, citings, index

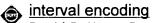
terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

4 XML query processing II: A comprehensive XQuery to SQL translation using dynamic





David DeHaan, David Toman, Mariano P. Consens, M. Tamer Özsu

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: T pdf(242.20 KB)

The W3C XQuery language recommendation, based on a hierarchical and ordered document model, supports a wide variety of constructs and use cases. There is a diversity of approaches and strategies for evaluating XQuery expressions, in many cases only dealing with limited subsets of the language. In this paper we describe an implementation approach that handles XQuery with arbitrarily-nested FLWR expressions, element constructors and built-in functions (including structural comparisons). Our propos ...

5 External memory algorithms and data structures: dealing with massive data



Jeffrey Scott Vitter

June 2001 ACM Computing Surveys (CSUR), Volume 33 Issue 2

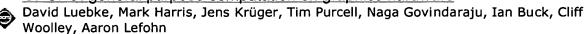
**Publisher: ACM Press** 

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: T pdf(828.46 KB) terms

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

6 GPGPU: general purpose computation on graphics hardware



August 2004 Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04

Publisher: ACM Press

Full text available: pdf(63.03 MB) Additional Information: full citation, abstract

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

7 Applications: Dynamic coordination of information management services for



processing dynamic web content

In-Young Ko, Ke-Thia Yao, Robert Neches

May 2002 Proceedings of the 11th international conference on World Wide Web

**Publisher: ACM Press** 

Full text available: pdf(1.15 MB)

Additional Information: full citation, abstract, references, citings, index terms

Dynamic Web content provides us with time-sensitive and continuously changing data. To glean up-to-date information, users need to regularly browse, collect and analyze this Web content. Without proper tool support this information management task is tedious, time-consuming and error prone, especially when the quantity of the dynamic Web content is large, when many information management services are needed to analyze it, and when underlying services/network are not completely reliable. This pap ...

Keywords: dynamic service coordination, dynamic web content, scalable componentbased software systems, semantic interoperability, web information management systems

Query optimization I: Access paths in the "Abe" statistical query facility



**Anthony Klug** 

June 1982 Proceedings of the 1982 ACM SIGMOD international conference on Management of data SIGMOD '82

Publisher: ACM Press

Full text available: pdf(1.07 MB)

Additional Information: full citation, abstract, references, citings

An increasingly important part of information processing today involves the taking of counts, sums, averages, and other statistical or aggregate quantities. The "Abe" query language is designed to make formulation of complicated aggregations simple. Access path selection in Abe finds efficient ways to execute these complicated queries. Access paths for Abe queries perform "aggregate joins", that is, they compute aggregate quantities at the same time as they join subqueries with parent queries. T ...

9 Level set and PDE methods for computer graphics



David Breen, Ron Fedkiw, Ken Museth, Stanley Osher, Guillermo Sapiro, Ross Whitaker August 2004 Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04

Publisher: ACM Press

Full text available: pdf(17.07 MB) Additional Information: full citation, abstract

Level set methods, an important class of partial differential equation (PDE) methods, define dynamic surfaces implicitly as the level set (iso-surface) of a sampled, evolving nD function. The course begins with preparatory material that introduces the concept of using partial differential equations to solve problems in computer graphics, geometric modeling and computer vision. This will include the structure and behavior of several different types of differential equations, e.g. the level set eq ...

10 Special issue on prototypes of deductive database systems: The aditi deductive database system



Jayen Vaghani, Kotagiri Ramamohanarao, David B. Kemp, Zoltan Somogyi, Peter J. Stuckey, Tim S. Leask, James Harland

April 1994 The VLDB Journal — The International Journal on Very Large Data Bases,

Volume 3 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: 📆 pdf(2.67 MB) Additional Information: full citation, abstract, references, citings

Deductive databases generalize relational databases by providing support for recursive views and non-atomic data. Aditi is a deductive system based on the client-server model; it is inherently multi-user and capable of exploiting parallelism on shared-memory multiprocessors. The back-end uses relational technology for efficiency in the management of disk-based data and uses optimization algorithms especially developed for the bottom-up evaluation of logical queries involving recursion. The front ...

Keywords: implementation, logic, multi-user, parallelism, relational database

11 Concurrent object-oriented programming

Gul Agha

September 1990 Communications of the ACM, Volume 33 Issue 9

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(2.06 MB) terms

Three significant trends have underscored the central role of concurrency in computing. First, there is increased use of interacting processes by individual users, for example, application programs running on X windows. Second, workstation networks have become a cost-effective mechanism for resource sharing and distributed problem solving. For example, loosely coupled problems, such as finding all the factors of large prime numbers, have been solved by utilizing ideal cycles on networks of ...

12 Computational strategies for object recognition

Paul Suetens, Pascal Fua, Andrew J. Hanson

March 1992 ACM Computing Surveys (CSUR), Volume 24 Issue 1

**Publisher: ACM Press** 

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(6.37 MB) terms, review

This article reviews the available methods for automated identification of objects in digital images. The techniques are classified into groups according to the nature of the computational strategy used. Four classes are proposed: (1) the simplest strategies, which work on data appropriate for feature vector classification, (2) methods that match models to symbolic data structures for situations involving reliable data and complex models, (3) approaches that fit models to the photometry and ...

Keywords: image understanding, model-based vision, object recognition

13 Graphical query interfaces for semistructured data: the QURSED system

Michalis Petropoulos, Yannis Papakonstantinou, Vasilis Vassalos

May 2005 ACM Transactions on Internet Technology (TOIT), Volume 5 Issue 2

Publisher: ACM Press

Full text available: pdf(6.88 MB) Additional Information: full citation, abstract, references, index terms

We describe the QURSED system for the declarative specification and automatic generation of Web-based query forms and reports (QFRs) for semistructured XML data. In OURSED, a OFR is formally described by its query set specification (QSS) which captures the complex query and reporting capabilities of the QFR and the associations of the query set specification with visual elements that implement these capabilities on a Web page.

The design-time component of QURSE ...

14 Concepts and paradigms of object-oriented programming

Peter Wegner

August 1990 ACM SIGPLAN OOPS Messenger, Volume 1 Issue 1

Publisher: ACM Press

Publisher: ACM Press

Full text available: pdf(5.52 MB) Additional Information: full citation, abstract, citings, index terms

We address the following questions for object-oriented programming: What is it? What are its goals? What are its origins? What are its paradigms? What are its design alternatives? What are its models of concurrency? What are its formal computational models? What comes after object-oriented programming? Starting from software engineering goals, we examine the origins and paradigms of object-oriented programming, explore its language design alternativ ...

15 XML parsing and stylesheets: Compiling XSLT 2.0 into XQuery 1.0

Achille Fokoue, Kristoffer Rose, Jérôme Siméon, Lionel Villard

May 2005 Proceedings of the 14th international conference on World Wide Web

Full text available: 🔁 pdf(143.18 KB) Additional Information: full citation, abstract, references, index terms

As XQuery is gathering momentum as the standard query language for XML, there is a growing interest in using it as an integral part of the XML application development infrastructure. In that context, one question which is often raised is how well XQuery interoperates with other XML languages, and notably with XSLT. XQuery 1.0 [16] and XSLT 2.0 [7] share a lot in common: they share XPath 2.0 as a common sub-language and have the same expressiveness. However, they are based on fairly different pro ...

Keywords: Web services, XML, XQuery, XSLT

16 WSQ/DSQ: a practical approach for combined querying of databases and the Web

Roy Goldman, Jennifer Widom

May 2000 ACM SIGMOD Record, Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00, Volume 29 Issue 2

Publisher: ACM Press

Full text available: pdf(223.65 KB)

Additional Information: full citation, abstract, references, citings, index terms

We present WSQ/DSQ (pronounced "wisk-disk"), a new approach for combining the query facilities of traditional databases with existing search engines on the Web. WSQ, for Web-Supported (Database) Queries, leverages results from Web searches to enhance SQL queries over a relational database. DSQ, for Database-Supported (Web) Queries, uses information stored in the database to enhance and explain Web searches. This paper focuses primarily on WSQ, describing a simple, lo ...

17 Research sessions: XML II: Approximate XML joins

Sudipto Guha, H. V. Jagadish, Nick Koudas, Divesh Srivastava, Ting Yu
June 2002 Proceedings of the 2002 ACM SIGMOD international conference on
Management of data SIGMOD '02

**Publisher: ACM Press** 

Full text available: pdf(1.25 MB)

Additional Information: full citation, abstract, references, citings, index terms

XML is widely recognized as the data interchange standard for tomorrow, because of its ability to represent data from a wide variety sources. Hence, XML is likely to be the format through which data from multiple sources is integrated. In this paper we study the



problem of integrating XML data sources through correlations realized as join operations. A challenging aspect of this operation is the XML document structure. Two documents might convey approximately or exactly the same information but m ...

18 Optimizing object queries using an effective calculus

Leonidas Fegaras, David Maier

December 2000 ACM Transactions on Database Systems (TODS), Volume 25 Issue 4

**Publisher: ACM Press** 

Full text available: pdf(641.65 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

Object-oriented databases (OODBs) provide powerful data abstractions and modeling facilities, but they generally lack a suitable framework for query processing and optimization. The development of an effective query optimizer is one of the key factors for OODB systems to successfully compete with relational systems, as well as to meet the performance requirements of many nontraditional applications. We propose an effective framework with a solid theoretical basis for optimizing OODB query I ...

**Keywords**: nested relations, object-oriented databases, query decorrelation, query optimization

19 <u>Model-based recognition in robot vision</u>



Roland T. Chin, Charles R. Dyer

March 1986 ACM Computing Surveys (CSUR), Volume 18 Issue 1

Publisher: ACM Press

Full text available: pdf(4.94 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

This paper presents a comparative study and survey of model-based object-recognition algorithms for robot vision. The goal of these algorithms is to recognize the identity, position, and orientation of randomly oriented industrial parts. In one form this is commonly referred to as the "bin-picking" problem, in which the parts to be recognized are presented in a jumbled bin. The paper is organized according to 2-D, 21/2-D, and 3-D object representations, which are used as the basis for ...

<sup>20</sup> Query optimization in the presence of limited access patterns



Daniela Florescu, Alon Levy, Ioana Manolescu, Dan Suciu

June 1999 ACM SIGMOD Record, Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99, Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(1.66 MB)

Additional Information: full citation, abstract, references, citings, index

We consider the problem of query optimization in the presence of limitations on access patterns to the data (i.e., when one must provide values for one of the attributes of a relation in order to obtain tuples). We show that in the presence of limited access patterns we must search a space of annotated query plans, where the annotations describe the inputs that must be given to the plan. We describe a theoretical and experimental analysis of the resulting search space and a ...

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

The ACM Digital Library Search: O The Guide

mining sort pattern templates

SEARCH



Feedback Report a problem Satisfaction survey

Terms used mining sort pattern templates

expanded form

Found 20,907 of 171,143

Sort results by

Display

results

 $\bigcirc$ relevance

Save results to a Binder Search Tips Open results in a new

Try an Advanced Search Try this search in The ACM Guide

window

Result page: 1 2 3 4 5 6 7 8 9 10

next

Best 200 shown

Results 1 - 20 of 200

Relevance scale

Data mining: Mining the peanut gallery: opinion extraction and semantic classification

of product reviews

Kushal Dave, Steve Lawrence, David M. Pennock

 $\nabla$ 

May 2003 Proceedings of the 12th international conference on World Wide Web

Publisher: ACM Press

Full text available: pdf(327.95 KB)

Additional Information: full citation, abstract, references, citings, index terms

The web contains a wealth of product reviews, but sifting through them is a daunting task. Ideally, an opinion mining tool would process a set of search results for a given item, generating a list of product attributes (quality, features, etc.) and aggregating opinions about each of them (poor, mixed, good). We begin by identifying the unique properties of this problem and develop a method for automatically distinguishing between positive and negative reviews. Our classifier draws on information ...

Keywords: document classification, opinion mining

2 Poster 3: content track: A probabilistic template-based approach to discovering



repetitive patterns in broadcast videos Peng Wang, Zhi-Qiang Liu, Shi-Qiang Yang

November 2005 Proceedings of the 13th annual ACM international conference on **Multimedia MULTIMEDIA '05** 

Publisher: ACM Press

Full text available: Top pdf(163.58 KB) Additional Information: full citation, abstract, references, index terms

There are usually repetitive sub-segments in broadcast videos, which may be associated with high-level concepts or events, e.g., news footage, repeated scores in basketball. Unsupervised mining techniques provide generic solutions to discovering such temporal patterns in various video genres, which are currently the subject of great interests to researchers working on multimedia content analysis. In this paper, we propose a novel approach to automatically detecting repetitive patterns in a video ...

**Keywords**: probabilistic template, repetitive pattern discovery, video mining

<u>Fast detection of communication patterns in distributed executions</u> Thomas Kunz, Michiel F. H. Seuren



### November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research

Publisher: IBM Press

Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

4 Mining semantics for large scale integration on the web: evidences, insights, and



<u>challenges</u>

Kevin Chen-Chuan Chang, Bin He, Zhen Zhang

December 2004 ACM SIGKDD Explorations Newsletter, Volume 6 Issue 2

Publisher: ACM Press

Full text available: pdf(466.64 KB) Additional Information: full citation, abstract, references

The Web has been rapidly "deepened" -- with myriad searchable databases online, where data are hidden behind query interfaces. Toward large scale integration over this "deep Web," we are facing a new challenge- With its dynamic and ad-hoc nature, such large scale integration mandates dynamic semantics discovery. That is, we must on-the-fly cope with "semantics" of dynamically discovered sources without pre-configured sourcespecific knowledge. To tackle this challenge, our initial ...

5 Bug localization: DynaMine: finding common error patterns by mining software revision histories





Benjamin Livshits, Thomas Zimmermann

September 2005 Proceedings of the 10th European software engineering conference held jointly with 13th ACM SIGSOFT international symposium on Foundations of software engineering ESEC/FSE-13

Publisher: ACM Press

Full text available: pdf(182.92 KB) Additional Information: full citation, abstract, references, index terms

A great deal of attention has lately been given to addressing software bugs such as errors in operating system drivers or security bugs. However, there are many other lesser known errors specific to individual applications or APIs and these violations of applicationspecific coding rules are responsible for a multitude of errors. In this paper we propose DynaMine, a tool that analyzes source code check-ins to find highly correlated method calls as well as common bug fixes in order to automatical ...

Keywords: coding patterns, data mining, dynamic analysis, error patterns, one-line check-ins, revision histories, software bugs

Exploration mining in diabetic patients databases: findings and conclusions



Wynne Hsu, Mong Li Lee, Bing Liu, Tok Wang Ling

August 2000 Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(136.53 KB) Additional Information: full citation, references, citings, index terms

7



Articles on microarray data mining: Towards interactive exploration of gene



### expression patterns

Daxin Jiang, Jian Pei, Aidong Zhang

December 2003 ACM SIGKDD Explorations Newsletter, Volume 5 Issue 2

Publisher: ACM Press

Full text available: Double Fig. 27.68 KB) Additional Information: Full citation, abstract, references

Analyzing coherent gene expression patterns is an important task in bioinformatics research and biomedical applications. Recently, various clustering methods have been adapted or proposed to identify clusters of co-expressed genes and recognize coherent expression patterns as the centroids of the clusters. However, the interpretation of coexpressed genes and coherent patterns mainly depends on the domain knowledge, which presents several challenges for coherent pattern mining and cannot be solv ...

8 Bug localization: PR-Miner: automatically extracting implicit programming rules and





detecting violations in large software code Zhenmin Li, Yuanyuan Zhou

September 2005 Proceedings of the 10th European software engineering conference held jointly with 13th ACM SIGSOFT international symposium on Foundations of software engineering ESEC/FSE-13

Publisher: ACM Press

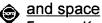
Full text available: Top pdf(228.58 KB) Additional Information: full citation, abstract, references, index terms

Programs usually follow many implicit programming rules, most of which are too tedious to be documented by programmers. When these rules are violated by programmers who are unaware of or forget about them, defects can be easily introduced. Therefore, it is highly desirable to have tools to automatically extract such rules and also to automatically detect violations. Previous work in this direction focuses on simple function-pair based programming rules and additionally requires programmer ...

**Keywords:** automated specification generation, automated violation detection, data mining for software engineering, pattern recognition, programming rules, static analysis

9 Poster papers: Finding surprising patterns in a time series database in linear time





Eamonn Keogh, Stefano Lonardi, Bill 'Yuan-chi' Chiu

July 2002 Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(686.15 KB)

Additional Information: full citation, abstract, references, citings, index terms

The problem of finding a specified pattern in a time series database (i.e. query by content) has received much attention and is now a relatively mature field. In contrast, the important problem of enumerating all surprising or interesting patterns has received far less attention. This problem requires a meaningful definition of "surprise", and an efficient search technique. All previous attempts at finding surprising patterns in time series use a very limited notion of surprise, and/or do not sc ...

**Keywords**: Markov Model, anomaly detection, feature extraction, novelty detection, suffix tree, time series

10 Background for association rules and cost estimate of selected mining algorithms Jia Liang Han, Ashley W. Plank





November 1996 Proceedings of the fifth international conference on Information and knowledge management

Publisher: ACM Press

Full text available: pdf(928.44 KB) Additional Information: full citation, references, citings, index terms

Keywords: algorithm analysis, association rules, background, data mining, memory management, optimization, scalability

11 Patterns and aspects: Detecting higher-level similarity patterns in programs

Hamid Abdul Basit, Stan Jarzabek

September 2005 Proceedings of the 10th European software engineering conference held jointly with 13th ACM SIGSOFT international symposium on Foundations of software engineering ESEC/FSE-13

Publisher: ACM Press

Full text available: pdf(641.23 KB) Additional Information: full citation, abstract, references, index terms

Cloning in software systems is known to create problems during software maintenance. Several techniques have been proposed to detect the same or similar code fragments in software, so-called simple clones. While the knowledge of simple clones is useful, detecting design-level similarities in software could ease maintenance even further, and also help us identify reuse opportunities. We observed that recurring patterns of simple clones - so-called structural clones - often indicate ...

**Keywords**: clone detection, similarity patterns, software clones

12 Industrial and practical experience track paper session 1: The volume and evolution



of web page templates

David Gibson, Kunal Punera, Andrew Tomkins

May 2005 Special interest tracks and posters of the 14th international conference on **World Wide Web** 

Publisher: ACM Press

Full text available: pdf(249.32 KB) Additional Information: full citation, abstract, references, index terms

Web pages contain a combination of unique content and template material, which is present across multiple pages and used primarily for formatting, navigation, and branding. We study the nature, evolution, and prevalence of these templates on the web. As part of this work, we develop new randomized algorithms for template extraction that perform approximately twenty times faster than existing approaches with similar quality. Our results show that 40--50% of the content on the web is templa ...

Keywords: algorithms, boilerplate, data cleaning, data mining, templates, web mining

13 Web and e-business application: Application run time estimation: a quality of service





metric for web-based data mining services

Shonali Krishnaswamy, Seng Wai Loke, Arkady Zaslavsky

March 2002 Proceedings of the 2002 ACM symposium on Applied computing

Publisher: ACM Press

Full text available: pdf(685.04 KB)

Additional Information: full citation, abstract, references, citings, index terms.

The emergence of Application Service Providers (ASP) hosting Internet-based data mining services is being seen as a viable alternative for organisations that value their knowledge resources but are constrained by the high cost of data mining software. Response time is an important Quality of Service (QoS) metric for web-based data mining service

providers. The ability to estimate the response time of data mining algorithms apriori benefits both clients and service providers. The advantage for th ...

Keywords: application run time estimation, data mining e-services, quality of service, rough sets

14 Emergent web patterns: The connectivity sonar: detecting site functionality by



structural patterns

Einat Amitay, David Carmel, Adam Darlow, Ronny Lempel, Aya Soffer

August 2003 Proceedings of the fourteenth ACM conference on Hypertext and hvpermedia

**Publisher: ACM Press** 

Full text available: pdf(153.40 KB)

Additional Information: full citation, abstract, references, citings, index terms.

Web sites today serve many different functions, such as corporate sites, search engines, e-stores, and so forth. As sites are created for different purposes, their structure and connectivity characteristics vary. However, this research argues that sites of similar role exhibit similar structural patterns, as the functionality of a site naturally induces a typical hyperlinked structure and typical connectivity patterns to and from the rest of the Web. Thus, the functionality of Web sites is refle ...

**Keywords**: link analysis, web IR, web graphs

15 Sequence mining in categorical domains: incorporating constraints





Mohammed J. Zaki

November 2000 Proceedings of the ninth international conference on Information and knowledge management

Publisher: ACM Press

Full text available: 📆 pdf(622.26 KB) Additional Information: full citation, references, citings, index terms

16 Computing curricula 2001



September 2001 Journal on Educational Resources in Computing (JERIC)

Publisher: ACM Press

Full text available: pdf(613.63 KB)

html(2.78 KB)

Additional Information: full citation, references, citings, index terms

17 Research track: Eliminating noisy information in Web pages for data mining



Lan Yi, Bing Liu, Xiaoli Li

August 2003 Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(468.87 KB)

Additional Information: full citation, abstract, references, citings, index terms

A commercial Web page typically contains many information blocks. Apart from the main content blocks, it usually has such blocks as navigation panels, copyright and privacy notices, and advertisements (for business purposes and for easy user access). We call these blocks that are not the main content blocks of the page the noisy blocks. We show that the information contained in these noisy blocks can seriously harm Web data mining. Eliminating these noises is thus of great importance. In this pa ...

**Keywords**: Web mining, noise detection, noise elimination

18 Web services and performance evaluation: Indexing web access-logs for pattern



<u>queries</u>

Alexandros Nanopoulos, Yannis Manolopoulos, Maciej Zakrzewicz, Tadeusz Morzy November 2002 Proceedings of the 4th international workshop on Web information and data management

**Publisher: ACM Press** 

Full text available: Topological pdf(187.24 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we develop a new indexing method for large web access-logs. We are concerned with pattern queries, which advocate the search for access sequences that contain certain query patterns. This kind of queries find applications in processing web-log mining results (e.g., finding typical/atypical access-sequences). The proposed method focuses on scalability to web-logs' sizes. For this reason, we examine the gains due to signature-trees, which can further improve the scalability to very ...

19 Knowledge discovery in data warehouses



Themistoklis Palpanas

September 2000 ACM SIGMOD Record, Volume 29 Issue 3

Publisher: ACM Press

Full text available: Ddf(240.77 KB) Additional Information: full citation, abstract, citings, index terms

As the size of data warehouses increase to several hundreds of gigabytes or terabytes, the need for methods and tools that will automate the process of knowledge extraction, or quide the user to subsets of the dataset that are of particular interest, is becoming prominent. In this survey paper we explore the problem of identifying and extracting interesting knowledge from large collections of data residing in data warehouses, by using data mining techniques. Such techniques have the ability to i ...

20 Sequence Mining: Efficient and robust feature extraction and pattern matching of time





series by a lattice structure

Polly Wan Po Man, Man Hon Wong

October 2001 Proceedings of the tenth international conference on Information and knowledge management

**Publisher: ACM Press** 

Full text available: pdf(1.48 MB)

Additional Information: full citation, abstract, references, citings, index terms

The efficiency of searching scaling-invariant and shifting-invariant shapes in a set of massive time series data can be improved if searching is performed on an approximated sequence which involves less data but contains all the significant features. However, commonly used smoothing techniques, such as moving averages and best-fitting polylines, usually miss important peaks and troughs and deform the time series. In addition, these techniques are not robust, as they often requires users to suppl ...

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: 

The ACM Digital Library O The Guide

anti monotonicity property mining templates

SEARCH



Feedback Report a problem Satisfaction survey

Terms used anti monotonicity property mining templates

 $\Box$ 

Found 11,946 of 171,143

Sort results by

Display

results

relevance  $\triangle$ 

expanded form

Save results to a Binder Search Tips

Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 200

window

Result page: 1 2 3 4 5 6 7 8 9 10

next Relevance scale

Best 200 shown

Exploratory mining and pruning optimizations of constrained associations rules

Raymond T. Ng, Laks V. S. Lakshmanan, Jiawei Han, Alex Pang

June 1998 ACM SIGMOD Record, Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98, Volume 27 Issue 2

**Publisher: ACM Press** 

Full text available: pdf(1.65 MB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

From the standpoint of supporting human-centered discovery of knowledge, the presentday model of mining association rules suffers from the following serious shortcomings: (i) lack of user exploration and control, (ii) lack of focus, and (iii) rigid notion of relationships. In effect, this model functions as a black-box, admitting little user interaction in between. We propose, in this paper, an architecture that opens up the black-box, and supports constraint-based, human-centered explorat ...

2 Multi Relational Data Mining (MRDM): Scalability and efficiency in multi-relational



data mining

Hendrik Blockeel, Michèle Sebag

July 2003 ACM SIGKDD Explorations Newsletter, Volume 5 Issue 1

**Publisher: ACM Press** 

Full text available: pdf(1.61 MB)

Additional Information: full citation, abstract, references, citings

Efficiency and Scalability have always been important concerns in the field of data mining, and are even more so in the multi-relational context, which is inherently more complex. The issue has been receiving an increasing amount of attention during the last few years, and quite a number of theoretical results, algorithms and implementations have been presented that explicitly aim at improving the efficiency and Scalability of multi-relational data mining approaches. With this article we attempt ...

3 Pattern discovery and forecasting: Mining sequential patterns with constraints in large



🙈 <u>data</u>bases

Jian Pei, Jiawei Han, Wei Wang

November 2002 Proceedings of the eleventh international conference on Information and knowledge management

Publisher: ACM Press

Full text available: pdf(252.19 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Constraints are essential for many sequential pattern mining applications. However, there

is no systematic study on constraint-based sequential pattern mining. In this paper, we investigate this issue and point out that the framework developed for constrained frequent-pattern mining does not fit our missions well. An extended framework is developed based on a sequential pattern growth methodology. Our study shows that constraints can be effectively and efficiently pushed deep into ...

4 A template model for multidimensional inter-transactional association rules Ling Feng, Jeffrey Xu Yu, Hongjun Lu, Jiawei Han



October 2002 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 11 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(313.62 KB) Additional Information: full citation, abstract, index terms

Multidimensional inter-transactional association rules extend the traditional association rules to describe more general associations among items with multiple properties across transactions. "After McDonald and Burger King open branches, KFC will open a branch two months later and one mile away" is an example of such rules. Since the number of potential inter-transactional association rules tends to be extremely large, mining intertransactional associations poses more challe ...

Keywords: Intra-transactional/inter-transactional association rules, Multidimensional context, Template model

5 Optimization of constrained frequent set queries with 2-variable constraints



Laks V. S. Lakshmanan, Raymond Ng, Jiawei Han, Alex Pang

June 1999 ACM SIGMOD Record, Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99, Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(1.65 MB)

Additional Information: full citation, abstract, references, citings, index terms

Currently, there is tremendous interest in providing ad-hoc mining capabilities in database management systems. As a first step towards this goal, in [15] we proposed an architecture for supporting constraint-based, human-centered, exploratory mining of various kinds of rules including associations, introduced the notion of constrained frequent set queries (CFQs), and developed effective pruning optimizations for CFQs with 1variable (1-var) constraints. While 1-var constraints a ...

6 Scalable data mining with model constraints



Minos Garofalakis, Rajeev Rastogi

December 2000 ACM SIGKDD Explorations Newsletter, Volume 2 Issue 2

Publisher: ACM Press

Full text available: 📆 pdf(1.15 MB) Additional Information: full citation, index terms

Keywords: data mining constraints, decision tables, sequential patterns

Mining the most interesting rules



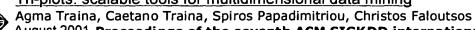
Roberto J. Bayardo, Rakesh Agrawal

August 1999 Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: 🔁 pdf(1.29 MB) Additional Information: full citation, references, citings, index terms

8 Tri-plots: scalable tools for multidimensional data mining



August 2001 Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(883.71 KB)

Additional Information: full citation, abstract, references, citings, index terms

We focus on the problem of finding patterns across two large, multidimensional datasets. For example, given feature vectors of healthy and of non-healthy patients, we want to answer the following questions: Are the two clouds of points separable? What is the smallest/largest pair-wise distance across the two datasets? Which of the two clouds does a new point (feature vector) come from? We propose a new tool, the tri-plot, and its generalization, the pq-plot, which help us answer the ...

9 Research track poster: Integration of profile hidden Markov model output into



Christopher Besemann, Anne Denton

August 2005 Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05

Publisher: ACM Press

Full text available: R pdf(794.15 KB) Additional Information: full citation, abstract, references, index terms

Scientific models typically depend on parameters. Preserving the parameter dependence of models in the pattern mining context opens up several applications. Within association rule mining (ARM), the choice of parameters can be studied with more flexibly then in traditional model building. Studying support, confidence, and other rule metrics as a function of model parameters allows conclusions on assumptions underlying the models. We present efficient techniques to handle multiple model output da ...

**Keywords:** association rule mining, model mining, profile hidden Markov model

Multi Relational Data Mining (MRDM): State of the art of graph-based data mining

Takashi Washio, Hiroshi Motoda

July 2003 ACM SIGKDD Explorations Newsletter, Volume 5 Issue 1

**Publisher: ACM Press** 

Full text available: pdf(1.20 MB) Additional Information: full citation, abstract, references, citings

The need for mining structured data has increased in the past few years. One of the best studied data structures in computer science and discrete mathematics are graphs. It can therefore be no surprise that graph based data mining has become quite popular in the last few years. This article introduces the theoretical basis of graph based data mining and surveys the state of the art of graph-based data mining. Brief descriptions of some representative approaches are provided as well.

Keywords: data mining, graph, graph-based data mining, path, structured data, tree

11 The segment support map: scalable mining of frequent itemsets

Laks V. S. Lakshmanan, Carson Kai-Sang Leung, Raymond T. Ng December 2000 ACM SIGKDD Explorations Newsletter, Volume 2 Issue 2

Publisher: ACM Press

Full text available: pdf(712.50 KB) Additional Information: full citation, citings, index terms

Keywords: association rule, frequent sets, scalable data mining

12 Research track paper: On mining cross-graph quasi-cliques

Jian Pei, Daxin Jiang, Aidong Zhang

August 2005 Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05

**Publisher: ACM Press** 

Full text available: 🔁 pdf(573.85 KB) Additional Information: full citation, abstract, references, index terms

Joint mining of multiple data sets can often discover interesting, novel, and reliable patterns which cannot be obtained solely from any single source. For example, in crossmarket customer segmentation, a group of customers who behave similarly in multiple markets should be considered as a more coherent and more reliable cluster than clusters found in a single market. As another example, in bioinformatics, by joint mining of gene expression data and protein interaction data, we can find cluster ...

**Keywords**: bioinformatics, graph mining, patterns

13 Research track posters: Generalizing the notion of support

Michael Steinbach, Pang-Ning Tan, Hui Xiong, Vipin Kumar

August 2004 Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining KDD '04

Publisher: ACM Press

Full text available: pdf(194.35 KB) Additional Information: full citation, abstract, references, index terms

The goal of this paper is to show that generalizing the notion of support can be useful in extending association analysis to non-traditional types of patterns and non-binary data. To that end, we describe a framework for generalizing support that is based on the simple, but useful observation that support can be viewed as the composition of two functions: a function that evaluates the strength or presence of a pattern in each object (transaction) and a function that summarizes these evaluations ...

**Keywords**: association analysis, hyperclique, support

14 A perspective on inductive databases

Luc De Raedt

December 2002 ACM SIGKDD Explorations Newsletter, Volume 4 Issue 2

**Publisher: ACM Press** 

Full text available: Topological pdf(179.64 KB) Additional Information: full citation, abstract, references, citings

Inductive databases tightly integrate databases with data mining. The key ideas are that data and patterns (or models) are handled in the same way and that an inductive query language allows the user to query and manipulate the patterns (or models) of interest. This paper proposes a simple and abstract model for inductive databases. We describe the basic formalism, a simple but fairly powerful inductive query language, some basics of reasoning for query optimization, and discuss some memory organ ...

**Keywords:** constraint-based mining, inductive databases, inductive querying

15 Scalable algorithms for mining large databases Rajeev Rastogi, Kyuseok Shim



### August 1999 Tutorial notes of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(4.11 MB) Additional Information: full citation, references, citings, index terms

16 Frequent patterns I: Selecting the right interestingness measure for association



patterns

Pang-Ning Tan, Vipin Kumar, Jaideep Srivastava

July 2002 Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(937.37 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

Many techniques for association rule mining and feature selection require a suitable metric to capture the dependencies among variables in a data set. For example, metrics such as support, confidence, lift, correlation, and collective strength are often used to determine the interestingness of association patterns. However, many such measures provide conflicting information about the interestingness of a pattern, and the best metric to use for a given application domain is rarely known. In this ...

**Keywords**: associations, contingency tables, interestingness measure

17 Logical foundations of object-oriented and frame-based languages



Michael Kifer, Georg Lausen, James Wu

July 1995 Journal of the ACM (JACM), Volume 42 Issue 4

Publisher: ACM Press

Full text available: pdf(7.52 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

We propose a novel formalism, called Frame Logic (abbr., F-logic), that accounts in a clean and declarative fashion for most of the structural aspects of object-oriented and frame-based languages. These features include object identity, complex objects, inheritance, polymorphic types, query methods, encapsulation, and others. In a sense, Flogic stands in the same relationship to the object-oriented paradigm as classical predicate calculus stands to relational programming. ...

**Keywords**: deductive databases, frame-based languages, logic programming, nonmonotonic inheritance, object-oriented programming, proof theory, semantics, typing

18 Exploratory mining via constrained frequent set queries



Raymond Ng, Laks V. S. Lakshmanan, Jiawei Han, Teresa Mah



Publisher: ACM Press

Full text available: pdf(353.58 KB)

Additional Information: full citation, abstract, references, citings, index

Although there have been many studies on data mining, to date there have been few research prototypes or commercial systems supporting comprehensive query-driven mining, which encourages interactive exploration of the data. Our thesis is that constraint constructs and the optimization they induce play a pivotal role in mining queries, thus substantially enhancing the usefulness and performance of the mining system. This is based on the analogy of declarative query languages like SQL and que ...

19 Multi Relational Data Mining (MRDM): Biological applications of multi-relational data



mining

David Page, Mark Craven

July 2003 ACM SIGKDD Explorations Newsletter, Volume 5 Issue 1

Publisher: ACM Press

Full text available: pdf(1.12 MB) Additional Information: full citation, abstract, references, citings

Biological databases contain a wide variety of data types, often with rich relational structure. Consequently multi-relational data mining techniques frequently are applied to biological data. This paper presents several applications of multi-relational data mining to biological data, taking care to cover a broad range of multi-relational data mining techniques.

20 Computing curricula 2001

September 2001 Journal on Educational Resources in Computing (JERIC)

**Publisher: ACM Press** 

Full text available: pdf(613.63 KB)

Additional Information: full citation, references, citings, index terms html(2.78 KB)

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player



Subscribe (Full Service) Register (Limited Service, Free) Login

O The Guide

anti monotonicity property mining templates





Feedback Report a problem Satisfaction survey

Terms used anti monotonicity property mining templates

Found 11,946 of 171,143

Sort results by Display

 $\Diamond$ relevance

Save results to a Binder Search Tips

Try an Advanced Search Try this search in The ACM Guide

expanded form Open results in a new results window

Results 1 - 20 of 200 Best 200 shown

Result page: 1 2 3 4 5 6 7 8 9 10 next

Relevance scale

Exploratory mining and pruning optimizations of constrained associations rules

Raymond T. Ng, Laks V. S. Lakshmanan, Jiawei Han, Alex Pang June 1998 ACM SIGMOD Record, Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98, Volume 27 Issue 2

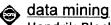
Publisher: ACM Press

Full text available: pdf(1.65 MB)

Additional Information: full citation, abstract, references, citings, index

From the standpoint of supporting human-centered discovery of knowledge, the presentday model of mining association rules suffers from the following serious shortcomings: (i) lack of user exploration and control, (ii) lack of focus, and (iii) rigid notion of relationships. In effect, this model functions as a black-box, admitting little user interaction in between. We propose, in this paper, an architecture that opens up the black-box, and supports constraint-based, human-centered explorat ...

2 Multi Relational Data Mining (MRDM): Scalability and efficiency in multi-relational



Hendrik Blockeel, Michèle Sebag

July 2003 ACM SIGKDD Explorations Newsletter, Volume 5 Issue 1

Publisher: ACM Press

Full text available: pdf(1.61 MB) Additional Information: full citation, abstract, references, citings

Efficiency and Scalability have always been important concerns in the field of data mining, and are even more so in the multi-relational context, which is inherently more complex. The issue has been receiving an increasing amount of attention during the last few years, and quite a number of theoretical results, algorithms and implementations have been presented that explicitly aim at improving the efficiency and Scalability of multi-relational data mining approaches. With this article we attempt ...

3 Pattern discovery and forecasting: Mining sequential patterns with constraints in large



databases

Jian Pei, Jiawei Han, Wei Wang

November 2002 Proceedings of the eleventh international conference on Information and knowledge management

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(252.19 KB)

Constraints are essential for many sequential pattern mining applications. However, there

is no systematic study on constraint-based sequential pattern mining. In this paper, we investigate this issue and point out that the framework developed for constrained frequent-pattern mining does not fit our missions well. An extended framework is developed based on a sequential pattern growth methodology. Our study shows that constraints can be effectively and efficiently pushed deep into ...

4 A template model for multidimensional inter-transactional association rules Ling Feng, Jeffrey Xu Yu, Hongjun Lu, Jiawei Han



October 2002 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 11 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: 📆 pdf(313.62 KB) Additional Information: full citation, abstract, index terms

Multidimensional inter-transactional association rules extend the traditional association rules to describe more general associations among items with multiple properties across transactions. "After McDonald and Burger King open branches, KFC will open a branch two months later and one mile away" is an example of such rules. Since the number of potential inter-transactional association rules tends to be extremely large, mining intertransactional associations poses more challe ...

Keywords: Intra-transactional/inter-transactional association rules, Multidimensional context, Template model

5 Optimization of constrained frequent set queries with 2-variable constraints



Laks V. S. Lakshmanan, Raymond Ng, Jiawei Han, Alex Pang

June 1999 ACM SIGMOD Record, Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99, Volume 28 Issue 2

**Publisher: ACM Press** 

Full text available: pdf(1.65 MB)

Additional Information: full citation, abstract, references, citings, index terms

Currently, there is tremendous interest in providing ad-hoc mining capabilities in database management systems. As a first step towards this goal, in [15] we proposed an architecture for supporting constraint-based, human-centered, exploratory mining of various kinds of rules including associations, introduced the notion of constrained frequent set queries (CFQs), and developed effective pruning optimizations for CFQs with 1variable (1-var) constraints. While 1-var constraints a ...

6 Scalable data mining with model constraints



Minos Garofalakis, Rajeev Rastogi

December 2000 ACM SIGKDD Explorations Newsletter, Volume 2 Issue 2

Publisher: ACM Press

Full text available: pdf(1.15 MB)

Additional Information: full citation, index terms

Keywords: data mining constraints, decision tables, sequential patterns

Mining the most interesting rules



Roberto J. Bayardo, Rakesh Agrawal

August 1999 Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(1.29 MB) Additional Information: full citation, references, citings, index terms

8 Tri-plots: scalable tools for multidimensional data mining

Agma Traina, Caetano Traina, Spiros Papadimitriou, Christos Faloutsos

August 2001 Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining

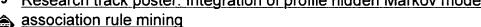
Publisher: ACM Press

Full text available: Ddf(883.71 KB)

Additional Information: full citation, abstract, references, citings, index terms

We focus on the problem of finding patterns across two large, multidimensional datasets. For example, given feature vectors of healthy and of non-healthy patients, we want to answer the following questions: Are the two clouds of points separable? What is the smallest/largest pair-wise distance across the two datasets? Which of the two clouds does a new point (feature vector) come from?We propose a new tool, the tri-plot, and its generalization, the pq-plot, which help us answer the ...

9 Research track poster: Integration of profile hidden Markov model output into



Christopher Besemann, Anne Denton

August 2005 Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05

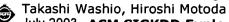
Publisher: ACM Press

Full text available: R pdf(794.15 KB) Additional Information: full citation, abstract, references, index terms

Scientific models typically depend on parameters. Preserving the parameter dependence of models in the pattern mining context opens up several applications. Within association rule mining (ARM), the choice of parameters can be studied with more flexibly then in traditional model building. Studying support, confidence, and other rule metrics as a function of model parameters allows conclusions on assumptions underlying the models. We present efficient techniques to handle multiple model output da ...

**Keywords:** association rule mining, model mining, profile hidden Markov model

Multi Relational Data Mining (MRDM): State of the art of graph-based data mining



July 2003 ACM SIGKDD Explorations Newsletter, Volume 5 Issue 1

Publisher: ACM Press

Full text available: pdf(1.20 MB) Additional Information: full citation, abstract, references, citings

The need for mining structured data has increased in the past few years. One of the best studied data structures in computer science and discrete mathematics are graphs. It can therefore be no surprise that graph based data mining has become quite popular in the last few years. This article introduces the theoretical basis of graph based data mining and surveys the state of the art of graph-based data mining. Brief descriptions of some representative approaches are provided as well.

Keywords: data mining, graph, graph-based data mining, path, structured data, tree

The segment support map: scalable mining of frequent itemsets

Laks V. S. Lakshmanan, Carson Kai-Sang Leung, Raymond T. Ng December 2000 ACM SIGKDD Explorations Newsletter, Volume 2 Issue 2

Publisher: ACM Press

Full text available: pdf(712.50 KB) Additional Information: full citation, citings, index terms



**Keywords**: association rule, frequent sets, scalable data mining

12 Research track paper: On mining cross-graph quasi-cliques

Jian Pei, Daxin Jiang, Aidong Zhang

August 2005 Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05

Publisher: ACM Press

Full text available: 📆 pdf(573.85 KB) Additional Information: full citation, abstract, references, index terms

Joint mining of multiple data sets can often discover interesting, novel, and reliable patterns which cannot be obtained solely from any single source. For example, in crossmarket customer segmentation, a group of customers who behave similarly in multiple markets should be considered as a more coherent and more reliable cluster than clusters found in a single market. As another example, in bioinformatics, by joint mining of gene expression data and protein interaction data, we can find cluster ...

**Keywords**: bioinformatics, graph mining, patterns

13 Research track posters: Generalizing the notion of support

Michael Steinbach, Pang-Ning Tan, Hui Xiong, Vipin Kumar

August 2004 Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining KDD '04

Publisher: ACM Press

Full text available: pdf(194.35 KB) Additional Information: full citation, abstract, references, index terms

The goal of this paper is to show that generalizing the notion of support can be useful in extending association analysis to non-traditional types of patterns and non-binary data. To that end, we describe a framework for generalizing support that is based on the simple, but useful observation that support can be viewed as the composition of two functions: a function that evaluates the strength or presence of a pattern in each object (transaction) and a function that summarizes these evaluations ...

**Keywords**: association analysis, hyperclique, support

14 A perspective on inductive databases

Luc De Raedt

December 2002 ACM SIGKDD Explorations Newsletter, Volume 4 Issue 2

Publisher: ACM Press

Full text available: pdf(179.64 KB) Additional Information: full citation, abstract, references, citings

Inductive databases tightly integrate databases with data mining. The key ideas are that data and patterns (or models) are handled in the same way and that an inductive query language allows the user to query and manipulate the patterns (or models) of interest. This paper proposes a simple and abstract model for inductive databases. We describe the basic formalism, a simple but fairly powerful inductive query language, some basics of reasoning for query optimization, and discuss some memory organ ...

**Keywords**: constraint-based mining, inductive databases, inductive querying

15 Scalable algorithms for mining large databases Rajeev Rastogi, Kyuseok Shim



### August 1999 Tutorial notes of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: 📆 pdf(4.11 MB) Additional Information: full citation, references, citings, index terms

16 Frequent patterns I: Selecting the right interestingness measure for association



patterns

Pang-Ning Tan, Vipin Kumar, Jaideep Srivastava

July 2002 Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining

**Publisher: ACM Press** 

Full text available: pdf(937.37 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

Many techniques for association rule mining and feature selection require a suitable metric to capture the dependencies among variables in a data set. For example, metrics such as support, confidence, lift, correlation, and collective strength are often used to determine the interestingness of association patterns. However, many such measures provide conflicting information about the interestingness of a pattern, and the best metric to use for a given application domain is rarely known. In this ...

Keywords: associations, contingency tables, interestingness measure

17 <u>Logical foundations of object-oriented and frame-based languages</u>



Michael Kifer, Georg Lausen, James Wu

July 1995 Journal of the ACM (JACM), Volume 42 Issue 4

Publisher: ACM Press

Full text available: pdf(7.52 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

We propose a novel formalism, called Frame Logic (abbr., F-logic), that accounts in a clean and declarative fashion for most of the structural aspects of object-oriented and frame-based languages. These features include object identity, complex objects, inheritance, polymorphic types, query methods, encapsulation, and others. In a sense, Flogic stands in the same relationship to the object-oriented paradigm as classical predicate calculus stands to relational programming. ...

**Keywords**: deductive databases, frame-based languages, logic programming, nonmonotonic inheritance, object-oriented programming, proof theory, semantics, typing

18 Exploratory mining via constrained frequent set queries



Raymond Ng, Laks V. S. Lakshmanan, Jiawei Han, Teresa Mah



Publisher: ACM Press

Full text available: pdf(353.58 KB)

Additional Information: full citation, abstract, references, citings, index

Although there have been many studies on data mining, to date there have been few research prototypes or commercial systems supporting comprehensive query-driven mining, which encourages interactive exploration of the data. Our thesis is that constraint constructs and the optimization they induce play a pivotal role in mining queries, thus substantially enhancing the usefulness and performance of the mining system. This is based on the analogy of declarative query languages like SQL and que ...

19 Multi Relational Data Mining (MRDM): Biological applications of multi-relational data



mining

David Page, Mark Craven

July 2003 ACM SIGKDD Explorations Newsletter, Volume 5 Issue 1

Publisher: ACM Press

Full text available: pdf(1.12 MB) Additional Information: full citation, abstract, references, citings

Biological databases contain a wide variety of data types, often with rich relational structure. Consequently multi-relational data mining techniques frequently are applied to biological data. This paper presents several applications of multi-relational data mining to biological data, taking care to cover a broad range of multi-relational data mining techniques.

20 Computing curricula 2001

September 2001 Journal on Educational Resources in Computing (JERIC)

Publisher: ACM Press

Full text available: pdf(613.63 KB)

Additional Information: full citation, references, citings, index terms html(2.78 KB)

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

The ACM Digital Library C The Guide Search:

anti monotonicity property

SEARCH



 $\nabla$ 

Feedback Report a problem Satisfaction survey

Terms used anti monotonicity property

Found 4,158 of 171,143

Sort results by

Display

results

1

relevance  $\bigcirc$ 

expanded form

Save results to a Binder Search Tips

Try an Advanced Search Try this search in The ACM Guide

next

Open results in a new window

Results 1 - 20 of 200

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

Relevance scale

Best 200 shown

Exploratory mining and pruning optimizations of constrained associations rules



Raymond T. Ng, Laks V. S. Lakshmanan, Jiawei Han, Alex Pang

June 1998 ACM SIGMOD Record, Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98, Volume 27 Issue 2

Publisher: ACM Press

Full text available: pdf(1.65 MB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

From the standpoint of supporting human-centered discovery of knowledge, the presentday model of mining association rules suffers from the following serious shortcomings: (i) lack of user exploration and control, (ii) lack of focus, and (iii) rigid notion of relationships. In effect, this model functions as a black-box, admitting little user interaction in between. We propose, in this paper, an architecture that opens up the black-box, and supports constraint-based, human-centered explorat ...

2 Pattern discovery and forecasting: Mining sequential patterns with constraints in large





databases

Jian Pei, Jiawei Han, Wei Wang

November 2002 Proceedings of the eleventh international conference on Information and knowledge management

Publisher: ACM Press

Full text available: pdf(252.19 KB)

Additional Information: full citation, abstract, references, citings, index terms

Constraints are essential for many sequential pattern mining applications. However, there is no systematic study on constraint-based sequential pattern mining. In this paper, we investigate this issue and point out that the framework developed for constrained frequent-pattern mining does not fit our missions well. An extended framework is developed based on a sequential pattern growth methodology. Our study shows that constraints can be effectively and efficiently pushed deep into ...

Optimization of constrained frequent set queries with 2-variable constraints



Laks V. S. Lakshmanan, Raymond Ng, Jiawei Han, Alex Pang

June 1999 ACM SIGMOD Record, Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99, Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(1.65 MB)

Additional Information: full citation, abstract, references, citings, index terms

Currently, there is tremendous interest in providing ad-hoc mining capabilities in database management systems. As a first step towards this goal, in [15] we proposed an architecture for supporting constraint-based, human-centered, exploratory mining of various kinds of rules including associations, introduced the notion of constrained frequent set queries (CFQs), and developed effective pruning optimizations for CFQs with 1variable (1-var) constraints. While 1-var constraints a ...

Routing: Metarouting

Timothy G. Griffin, João Luís Sobrinho

August 2005 Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '05

Publisher: ACM Press

Full text available: 🔁 pdf(258.67 KB) Additional Information: full citation, abstract, references, index terms

There is a shortage of routing protocols that meet the needs of network engineers. This has led to BGP being pressed into service as an IGP, despite its lack of convergence guarantees. The development, standardization, and deployment of routing protocols, or even minor changes to existing protocols, are very difficult tasks. We present an approach called Metarouting that defines routing protocols using a high-level and declarative language. Once an interpreter for a metarouting language is imple ...

**Keywords**: algebraic routing, path algebras, routing protocols

5 Logical foundations of object-oriented and frame-based languages

Michael Kifer, Georg Lausen, James Wu

July 1995 Journal of the ACM (JACM), Volume 42 Issue 4

Publisher: ACM Press

Full text available: pdf(7.52 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, citings, index

terms, review

We propose a novel formalism, called Frame Logic (abbr., F-logic), that accounts in a clean and declarative fashion for most of the structural aspects of object-oriented and frame-based languages. These features include object identity, complex objects, inheritance, polymorphic types, query methods, encapsulation, and others. In a sense, Flogic stands in the same relationship to the object-oriented paradigm as classical predicate calculus stands to relational programming. ...

Keywords: deductive databases, frame-based languages, logic programming, nonmonotonic inheritance, object-oriented programming, proof theory, semantics, typing

Scalable data mining with model constraints

Minos Garofalakis, Rajeev Rastogi

December 2000 ACM SIGKDD Explorations Newsletter, Volume 2 Issue 2

Publisher: ACM Press

Full text available: pdf(1.15 MB) Additional Information: full citation, index terms

Keywords: data mining constraints, decision tables, sequential patterns

7 Mining the most interesting rules

Roberto J. Bayardo, Rakesh Agrawal August 1999 Proceedings of the fifth ACM SIGKDD international conference on

#### Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(1.29 MB) Additional Information: full citation, references, citings, index terms

A perspective on inductive databases

Luc De Raedt

December 2002 ACM SIGKDD Explorations Newsletter, Volume 4 Issue 2

Publisher: ACM Press

Full text available: pdf(179.64 KB) Additional Information: full citation, abstract, references, citings

Inductive databases tightly integrate databases with data mining. The key ideas are that data and patterns (or models) are handled in the same way and that an inductive query language allows the user to query and manipulate the patterns (or models) of interest. This paper proposes a simple and abstract model for inductive databases. We describe the basic formalism, a simple but fairly powerful inductive query language, some basics of reasoning for query optimization, and discuss some memory organ ...

**Keywords:** constraint-based mining, inductive databases, inductive querying

9 Frequent patterns I: Selecting the right interestingness measure for association



patterns

Pang-Ning Tan, Vipin Kumar, Jaideep Srivastava

July 2002 Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(937.37 KB) terms, review

Many techniques for association rule mining and feature selection require a suitable metric to capture the dependencies among variables in a data set. For example, metrics such as support, confidence, lift, correlation, and collective strength are often used to determine the interestingness of association patterns. However, many such measures provide conflicting information about the interestingness of a pattern, and the best metric to use for a given application domain is rarely known. In this ...

**Keywords**: associations, contingency tables, interestingness measure

10 Tri-plots: scalable tools for multidimensional data mining

Agma Traina, Caetano Traina, Spiros Papadimitriou, Christos Faloutsos

August 2001 Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(883.71 KB) terms

We focus on the problem of finding patterns across two large, multidimensional datasets. For example, given feature vectors of healthy and of non-healthy patients, we want to answer the following questions: Are the two clouds of points separable? What is the smallest/largest pair-wise distance across the two datasets? Which of the two clouds does a new point (feature vector) come from? We propose a new tool, the tri-plot, and its generalization, the pq-plot, which help us answer the ...

Subtyping recursive types



Roberto M. Amadio, Luca Cardelli

September 1993 ACM Transactions on Programming Languages and Systems (TOPLAS), Volume 15 Issue 4

**Publisher: ACM Press** 

Full text available: pdf(3.29 MB)

Additional Information: full citation, abstract, references, citings, index terms

We investigate the interactions of subtyping and recursive types, in a simply typed &lgr;calculus. The two fundamental questions here are whether two (recursive)types are in the subtype relation and whether a term has a type. To address the first question, we relate various definitions of type equivalence and subtyping that are induced by a model, an ordering on infinite trees, an algorithm, and a set of type rules. We show soundness and completeness among the rules, the algorithm, and the ...

Keywords: coercions, lambda-calculus, partial-equivalence relations, recursive types, regular trees, subtyping, tree orderings, type equivalence, typechecking algorithm

12 Searching for dependencies at multiple abstraction levels



Toon Calders, Raymond T. Ng, Jef Wijsen

September 2002 ACM Transactions on Database Systems (TODS), Volume 27 Issue 3

Publisher: ACM Press

Full text available: R pdf(411.24 KB)

Additional Information: full citation, abstract, references, citings, index

The notion of roll-up dependency (RUD) extends functional dependencies with generalization hierarchies. RUDs can be applied in OLAP and database design. The problem of discovering RUDs in large databases is at the center of this paper. An algorithm is provided that relies on a number of theoretical results. The algorithm has been implemented; results on two real-life datasets are given. The extension of functional dependency (FD) with roll-ups turns out to capture meaningful rules that are outsi ...

Keywords: Data mining, functional dependencies, knowledge discovery, online analytical processing

13 Ordered chaining calculi for first-order theories of transitive relations



Leo Bachmair, Harald Ganzinger

November 1998 Journal of the ACM (JACM), Volume 45 Issue 6

Publisher: ACM Press

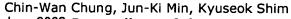
Full text available: pdf(279.40 KB)

Additional Information: full citation, abstract, references, citings, index

We propose inference systems for binary relations that satisfy composition laws such as transitivity. Our inference mechanisms are based on standard techniques from term rewriting and represent a refinement of chaining methods as they are used in the context of resolution-type theorem proving. We establish the refutational completeness of these calculi and prove that our methods are compatible with the usual simplification techniques employed in refutational theorem provers, such as subsump ...

Keywords: chaining calculi, equational logic, reduction orderings, rewrite systems, term rewriting, transitive relations

Research sessions: path indexing: APEX: an adaptive path index for XML data



June 2002 Proceedings of the 2002 ACM SIGMOD international conference on



#### Management of data SIGMOD '02

**Publisher: ACM Press** 

Full text available: pdf(1.16 MB)

Additional Information: full citation, abstract, references, citings, index terms

The emergence of the Web has increased interests in XML data. XML query languages such as XQuery and XPath use label paths to traverse the irregularly structured data. Without a structural summary and efficient indexes, query processing can be quite inefficient due to an exhaustive traversal on XML data. To overcome the inefficiency, several path indexes have been proposed in the research community. Traditional indexes generally record all label paths from the root element in XML data. Such path ...

15 Research track posters: Generalizing the notion of support

Michael Steinbach, Pang-Ning Tan, Hui Xiong, Vipin Kumar

August 2004 Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining KDD '04

Publisher: ACM Press

Full text available: pdf(194.35 KB) Additional Information: full citation, abstract, references, index terms

The goal of this paper is to show that generalizing the notion of support can be useful in extending association analysis to non-traditional types of patterns and non-binary data. To that end, we describe a framework for generalizing support that is based on the simple, but useful observation that support can be viewed as the composition of two functions: a function that evaluates the strength or presence of a pattern in each object (transaction) and a function that summarizes these evaluations ...

**Keywords**: association analysis, hyperclique, support

16 Exploratory mining via constrained frequent set queries

Raymond Ng, Laks V. S. Lakshmanan, Jiawei Han, Teresa Mah

June 1999 ACM SIGMOD Record, Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99, Volume 28 Issue 2

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(353.58 KB) terms

Although there have been many studies on data mining, to date there have been few research prototypes or commercial systems supporting comprehensive query-driven mining, which encourages interactive exploration of the data. Our thesis is that constraint constructs and the optimization they induce play a pivotal role in mining queries, thus substantially enhancing the usefulness and performance of the mining system. This is based on the analogy of declarative query languages like SQL and que ...

17 Research track poster: Integration of profile hidden Markov model output into

association rule mining

Christopher Besemann, Anne Denton

August 2005 Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05

Publisher: ACM Press

Full text available: Top pdf(794.15 KB) Additional Information: full citation, abstract, references, index terms

Scientific models typically depend on parameters. Preserving the parameter dependence of models in the pattern mining context opens up several applications. Within association rule mining (ARM), the choice of parameters can be studied with more flexibly then in traditional model building. Studying support, confidence, and other rule metrics as a function of model parameters allows conclusions on assumptions underlying the models. We present efficient techniques to handle multiple model output da ...

Keywords: association rule mining, model mining, profile hidden Markov model

18 Proving the correctness of reactive systems using sized types John Hughes, Lars Pareto, Amr Sabry



January 1996 Proceedings of the 23rd ACM SIGPLAN-SIGACT symposium on Principles of programming languages POPL '96

Publisher: ACM Press

Full text available: 🔁 pdf(1.47 MB)

Additional Information: full citation, references, citings, index terms

19 Multi Relational Data Mining (MRDM): Scalability and efficiency in multi-relational





data mining

Hendrik Blockeel, Michèle Sebag

July 2003 ACM SIGKDD Explorations Newsletter, Volume 5 Issue 1

**Publisher: ACM Press** 

Full text available: pdf(1.61 MB)

Additional Information: full citation, abstract, references, citings

Efficiency and Scalability have always been important concerns in the field of data mining, and are even more so in the multi-relational context, which is inherently more complex. The issue has been receiving an increasing amount of attention during the last few years, and quite a number of theoretical results, algorithms and implementations have been presented that explicitly aim at improving the efficiency and Scalability of multi-relational data mining approaches. With this article we attempt ...

<sup>20</sup> Scalable algorithms for mining large databases





Rajeev Rastogi, Kyuseok Shim

August 1999 Tutorial notes of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining

Publisher: ACM Press

Full text available: pdf(4.11 MB)

Additional Information: full citation, references, citings, index terms

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2006 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player